

Tru64 UNIX

Kernel Debugging

Part Number: AA-RH99A-TE

July 1999

Product Version: Tru64 UNIX Version 5.0 or higher

This manual explains how to use tools to debug a kernel and analyze a crash dump of the Tru64 UNIX (formerly DIGITAL UNIX) operating system. Also, this manual explains how to write extensions to the kernel debugging tools.

© 1999 Compaq Computer Corporation

COMPAQ and the Compaq logo Registered in U.S. Patent and Trademark Office. Alpha and Tru64 are trademarks of Compaq Information Technologies Group, L.P in the United States and other countries.

Microsoft and Windows are trademarks of Microsoft Corporation in the United States and other countries. UNIX and The Open Group are trademarks of The Open Group. All other product names mentioned herein may be trademarks of their respective companies.

Confidential computer software. Valid license from Compaq required for possession, use, or copying. Consistent with FAR 12.211 and 12.212, Commercial Computer Software, Computer Software Documentation, and Technical Data for Commercial Items are licensed to the U.S. Government under vendor's standard commercial license.

Compaq shall not be liable for technical or editorial errors or omissions contained herein. The information in this document is provided "as is" without warranty of any kind and is subject to change without notice. The warranties for Compaq products are set forth in the express limited warranty statements accompanying such products. Nothing herein should be construed as constituting an additional warranty.

Contents

About This Manual

1 Introduction to Kernel Debugging

1.1	Linking a Kernel Image for Debugging	1-1
1.2	Debugging Kernel Programs	1-3
1.3	Debugging the Running Kernel	1-3
1.4	Analyzing a Crash Dump File	1-5

2 Kernel Debugging Utilities

2.1	The dbx Debugger	2-2
2.1.1	Invoking the dbx Debugger for Kernel Debugging	2-2
2.1.2	Debugging Stripped Images	2-3
2.1.3	Specifying the Location of Loadable Modules for Crash Dumps	2-4
2.1.4	Examining Memory Contents	2-5
2.1.5	Printing the Values of Variables and Data Structures	2-6
2.1.6	Displaying a Data Structure Format	2-6
2.1.7	Debugging Multiple Threads	2-7
2.1.8	Examining the Exception Frame	2-7
2.1.9	Examining the User Program Stack	2-8
2.1.10	Extracting the Preserved Message Buffer	2-10
2.1.11	Debugging on SMP Systems	2-10
2.2	The kdbx Debugger	2-12
2.2.1	Beginning a kdbx Session	2-12
2.2.2	The kdbx Debugger Commands	2-13
2.2.3	Using kdbx Debugger Extensions	2-15
2.2.3.1	Displaying the Address Resolution Protocol Table	2-16
2.2.3.2	Performing Commands on Array Elements	2-16
2.2.3.3	Displaying the Buffer Table	2-18
2.2.3.4	Displaying the Callout Table and Absolute Callout Table	2-18
2.2.3.5	Casting Information Stored in a Specific Address	2-19
2.2.3.6	Displaying Machine Configuration	2-19
2.2.3.7	Converting the Base of Numbers	2-20
2.2.3.8	Displaying CPU Use Statistics	2-20

2.2.3.9	Disassembling Instructions	2-21
2.2.3.10	Displaying Remote Exported Entries	2-21
2.2.3.11	Displaying the File Table	2-21
2.2.3.12	Displaying the udb and tcb Tables	2-22
2.2.3.13	Performing Commands on Lists	2-22
2.2.3.14	Displaying the lockstats Structures	2-24
2.2.3.15	Displaying lockinfo Structures	2-25
2.2.3.16	Displaying the Mount Table	2-26
2.2.3.17	Displaying the Namecache Structures	2-27
2.2.3.18	Displaying Processes' Open Files	2-27
2.2.3.19	Converting the Contents of Memory to Symbols	2-28
2.2.3.20	Displaying the Process Control Block for a Thread	2-28
2.2.3.21	Formatting Command Arguments	2-28
2.2.3.22	Displaying the Process Table	2-29
2.2.3.23	Converting an Address to a Procedure name	2-30
2.2.3.24	Displaying Sockets from the File Table	2-30
2.2.3.25	Displaying a Summary of the System Information	2-30
2.2.3.26	Displaying a Summary of Swap Space	2-31
2.2.3.27	Displaying the Task Table	2-31
2.2.3.28	Displaying Information About Threads	2-32
2.2.3.29	Displaying a Stack Trace of Threads	2-32
2.2.3.30	Displaying a u Structure	2-33
2.2.3.31	Displaying References to the ucred Structure	2-34
2.2.3.32	Removing Aliases	2-36
2.2.3.33	Displaying the vnode Table	2-36
2.3	The kdebug Debugger	2-37
2.3.1	Getting Ready to Use the kdebug Debugger	2-39
2.3.2	Invoking the kdebug Debugger	2-41
2.3.3	Diagnosing kdebug Setup Problems	2-42
2.3.4	Notes on Using the kdebug Debugger	2-44
2.4	The crashdc Utility	2-44

3 Writing Extensions to the kdbx Debugger

3.1	Basic Considerations for Writing Extensions	3-1
3.2	Standard kdbx Library Functions	3-2
3.2.1	Special kdbx Extension Data Types	3-2
3.2.2	Converting an Address to a Procedure Name	3-3
3.2.3	Getting a Representation of an Array Element	3-4
3.2.4	Retrieving an Array Element Value	3-4
3.2.5	Returning the Size of an Array	3-6
3.2.6	Casting a Pointer to a Data Structure	3-6

3.2.7	Checking Arguments Passed to an Extension	3-7
3.2.8	Checking the Fields in a Structure	3-7
3.2.9	Setting the kdbx Context	3-8
3.2.10	Passing Commands to the dbx Debugger	3-9
3.2.11	Dereferencing a Pointer	3-9
3.2.12	Displaying the Error Messages Stored in Fields	3-10
3.2.13	Converting a Long Address to a String Address	3-10
3.2.14	Freeing Memory	3-11
3.2.15	Passing Commands to the kdbx Debugger	3-11
3.2.16	Getting the Address of an Item in a Linked List	3-13
3.2.17	Passing an Extension to kdbx	3-14
3.2.18	Getting the Next Token as an Integer	3-14
3.2.19	Getting the Next Token as a String	3-15
3.2.20	Displaying a Message	3-16
3.2.21	Displaying Status Messages	3-16
3.2.22	Exiting from an Extension	3-17
3.2.23	Reading the Values in Structure Fields	3-17
3.2.24	Returning a Line of kdbx Output	3-18
3.2.25	Reading an Area of Memory	3-18
3.2.26	Reading the Response to a kdbx Command	3-19
3.2.27	Reading Symbol Representations	3-20
3.2.28	Reading a Symbol's Address	3-20
3.2.29	Reading the Value of a Symbol	3-21
3.2.30	Getting the Address of a Data Representation	3-21
3.2.31	Converting a String to a Number	3-22
3.3	Examples of kdbx Extensions	3-22
3.4	Compiling Custom Extensions	3-35
3.5	Debugging Custom Extensions	3-36

4 Crash Analysis Examples

4.1	Guidelines for Examining Crash Dump Files	4-1
4.2	Identifying a Crash Caused by a Software Problem	4-2
4.2.1	Using dbx to Determine the Cause of a Software Panic ...	4-2
4.2.2	Using kdbx to Determine the Cause of a Software Panic ..	4-3
4.3	Identifying a Hardware Exception	4-4
4.3.1	Using dbx to Determine the Cause of a Hardware Error ..	4-4
4.3.2	Using kdbx to Determine the Cause of a Hardware Error	4-7
4.4	Finding a Panic String in a Thread Other Than the Current Thread	4-8
4.5	Identifying the Cause of a Crash on an SMP System	4-9

A Output from the crashdc Command

Index

Examples

3-1	Template Extension Using Lists	3-23
3-2	Extension That Uses Linked Lists: callout.c	3-24
3-3	Template Extensions Using Arrays	3-27
3-4	Extension That Uses Arrays: file.c	3-28
3-5	Extension That Uses Global Symbols: sum.c	3-34

Figures

2-1	Using a Gateway System During Remote Debugging	2-38
-----	--	------

Tables

2-1	The dbx Address Modes	2-5
-----	-----------------------------	-----

About This Manual

This manual provides information on the tools used to debug a kernel and analyze a crash dump file of the Tru64™ UNIX (formerly DIGITAL UNIX) operating system. It also explains how to write extensions to the kernel debugging tools. You can use extensions to display customized information from kernel data structures or a crash dump file.

Audience

This manual is intended for system programmers who write programs that use kernel data structures and are built into the kernel. It is also intended for system administrators who are responsible for managing the operating system. System programmers and administrators should have in-depth knowledge of operating system concepts, commands, and utilities.

New and Changed Features

The following list describes changes that have been made to this manual for Tru64 UNIX Version 5.0:

- The former Chapter 4, *Managing Crash Dumps*, has been deleted and its contents have been moved to the *System Administration* manual. All information on that subject is now in one manual. The *System Administration* manual was chosen because many aspects of managing crash dumps (such as storage considerations and default settings) are handled by a system administrator, often during system installation.
- Crash dumps are now compressed by default and are stored in compressed crash dump files. These are named `vmzcore.n` to differentiate them from the uncompressed `vmcore.n` files. Starting with Version 5.0, all the Tru64 UNIX debugging tools can read `vmzcore.n` as well as `vmcore.n` files. Examples throughout this manual have been updated to show use of `vmzcore.n` files.
- When debugging a crash dump with `dbx` or `kdbx`, you can examine the call stack of the user program whose execution precipitated the kernel crash. For more information, see *Section 2.1.9*.
- If a loadable kernel module was moved to another location after a kernel crash, you can specify the directory path where `dbx` should look for the module. For more information, see *Section 2.1.3*.

Organization

This manual consists of four chapters and one appendix:

<i>Chapter 1</i>	Introduces the concepts of kernel debugging and crash dump analysis.
<i>Chapter 2</i>	Describes the tools used to debug kernels and analyze crash dump files.
<i>Chapter 3</i>	Describes how to write a <code>kdbx</code> debugger extension. This chapter assumes you have purchased and installed a Tru64 UNIX Source Kit and so have access to source files.
<i>Chapter 4</i>	Provides background information useful for and examples of analyzing crash dump files.
<i>Appendix A</i>	Contains example output from the <code>crashdc</code> utility.

Related Documents

For additional information, refer to the following manuals:

- The *Alpha Architecture Reference Manual* describes how the operating system interfaces with the Alpha hardware.
- The *Alpha Architecture Handbook* gives an overview of the Alpha hardware architecture and describes the 64-bit Alpha RISC (Reduced Instruction Set Computing) instruction set.
- The *Installation Guide* and *Installation Guide — Advanced Topics* describe how to install your operating system.
- The *System Administration* manual provides information on managing and monitoring your system, including managing crash dumps.
- The *Programmer's Guide* provides information on the tools, specifically the `dbx` debugger, for programming on the Tru64 UNIX operating system. This manual also provides information about creating configurable kernel subsystems.
- The *Writing Kernel Modules* manual discusses how to code kernel modules (single binary images) that can be statically loaded as part of the `/vmunix` kernel or dynamically loaded into memory, that enhance the functionality of the Unix kernel.

Icons on Tru64 UNIX Printed Manuals

The printed version of the Tru64 UNIX documentation uses letter icons on the spines of the manuals to help specific audiences quickly find the manuals that meet their needs. (You can order the printed documentation from Compaq.) The following list describes this convention:

- G Manuals for general users
- S Manuals for system and network administrators
- P Manuals for programmers
- R Manuals for reference page users

Some manuals in the documentation help meet the needs of several audiences. For example, the information in some system manuals is also used by programmers. Keep this in mind when searching for information on specific topics.

The *Documentation Overview* provides information on all of the manuals in the Tru64 UNIX documentation set.

Reader's Comments

Compaq welcomes any comments and suggestions you have on this and other Tru64 UNIX manuals.

You can send your comments in the following ways:

- Fax: 603-884-0120 Attn: UBPG Publications, ZKO3-3/Y32
- Internet electronic mail: `readers_comment@zk3.dec.com`

A Reader's Comment form is located on your system in the following location:

```
/usr/doc/readers_comment.txt
```

Please include the following information along with your comments:

- The full title of the manual and the order number. (The order number appears on the title page of printed and PDF versions of a manual.)
- The section numbers and page numbers of the information on which you are commenting.
- The version of Tru64 UNIX that you are using.
- If known, the type of processor that is running the Tru64 UNIX software.

The Tru64 UNIX Publications group cannot respond to system problems or technical support inquiries. Please address technical questions to your local system vendor or to the appropriate Compaq technical support office. Information provided with the software media explains how to send problem reports to Compaq.

Conventions

The following conventions are used in this manual:

%	
\$	A percent sign represents the C shell system prompt. A dollar sign represents the system prompt for the Bourne, Korn, and POSIX shells.
#	A number sign represents the superuser prompt.
% cat	Boldface type in interactive examples indicates typed user input.
<i>file</i>	Italic (slanted) type indicates variable values, placeholders, and function argument names.
[]	
{ }	In syntax definitions, brackets indicate items that are optional and braces indicate items that are required. Vertical bars separating items inside brackets or braces indicate that you choose one item from among those listed.
:	
:	A vertical ellipsis indicates that a portion of an example that would normally be present is not shown.
cat(1)	A cross-reference to a reference page includes the appropriate section number in parentheses. For example, <code>cat(1)</code> indicates that you can find information on the <code>cat</code> command in Section 1 of the reference pages.
Ctrl/ <i>x</i>	This symbol indicates that you hold down the first named key while pressing the key or mouse button that follows the slash. In examples, this key combination is enclosed in a box (for example, Ctrl/C).

Introduction to Kernel Debugging

Kernel debugging is a task normally performed by systems engineers writing kernel programs. A kernel program is one that is built as part of the kernel and that references kernel data structures. System administrators might also debug the kernel in the following situations:

- A process is hung or stops running unexpectedly
- The need arises to examine, and possibly modify, kernel parameters
- The system itself hangs, panics, or crashes

This manual describes how to debug kernel programs and the kernel. It also includes information about analyzing crash dump files.

In addition to the information provided here, tracing a kernel problem can require a basic understanding of one or more of the following technical areas:

- The hardware architecture

See the *Alpha Architecture Handbook* for an overview of the Alpha hardware architecture and a description of the 64-bit Alpha RISC instruction set.

- The internal design of the operating system at a source code and data structure level

See the *Alpha Architecture Reference Manual* for information on how the Tru64 UNIX operating system interfaces with the hardware.

This chapter provides an overview of the following topics:

- Linking a kernel image prior to debugging for systems that are running a kernel built at boot time. (Section 1.1)
- Debugging kernel programs (Section 1.2)
- Debugging the running kernel (Section 1.3)
- Analyzing a crash dump file (Section 1.4)

1.1 Linking a Kernel Image for Debugging

By default, the kernel is a statically linked image that resides in the file `/vmunix`. However, your system might be configured so that it is linked at bootstrap time. Rather than being a bootable image, the boot file is a

text file that describes the hardware and software that will be present on the running system. Using this information, the bootstrap linker links the modules that are needed to support this hardware and software. The linker builds the kernel directly into memory.

You cannot directly debug a bootstrap-linked kernel because you must supply the name of an image to the kernel debugging tools. Without the image, the tools have no access to symbol names, variable names, and so on. Therefore, the first step in any kernel debugging effort is to determine whether your kernel was linked at bootstrap time. If the kernel was linked at bootstrap time, you must then build a kernel image file to use for debugging purposes.

The best way to determine whether your system is bootstrap linked or statically linked is to use the `file` command to test the type of file from which your system was booted. If your system is a bootstrap-linked system, it was booted from an ASCII text file; otherwise, it was booted from an executable image file. For example, issue the following command to determine the type of file from which your system was booted:

```
#!/usr/bin/file `usr/sbin/sizer -b`  
/etc/sysconfigtab: ascii text
```

The `sizer -b` command returns the name of the file from which the system was booted. This file name is input to the `file` command, which determines that the system was booted from an ASCII text file. The output shown in the preceding example indicates that the system is a bootstrap-linked system. If the system had been booted from an executable image file named `vmunix`, the output from the `file` command would have appeared as follows:

```
vmunix:COFF format alpha executable or object module  
not stripped
```

If your system is running a bootstrap-linked kernel, build a kernel image that is identical to the bootstrap-linked kernel your system is running, by entering the following command:

```
# /usr/bin/ld -o vmunix.image `usr/sbin/sizer -m`
```

The output from the `sizer -m` command is a list of the exact modules and linker flags used to build the currently running bootstrap-linked kernel. This output causes the `ld` command to create a kernel image that is identical to the bootstrap-linked kernel running on your system. The kernel image is written to the file named by the `-o` flag, in this case the `vmunix.image` file.

Once you create this image, you can debug the kernel as described in this manual, using the `dbx`, `kdbx`, and `kdebug` debuggers. When you invoke the `dbx` or `kdbx` debugger, remember to specify the name of the kernel image file you created with the `ld` command, such as the `vmunix.image` file shown here.

When you are finished debugging the kernel, you can remove the kernel image file you created for debugging purposes.

1.2 Debugging Kernel Programs

Kernel programs can be difficult to debug because you normally cannot control kernel execution. To make debugging kernel programs more convenient, the system provides the `kdebug` debugger. The `kdebug` debugger is code that resides inside the kernel and allows you to use the `dbx` debugger to control execution of a running kernel in the same manner as you control execution of a user space program. To debug a kernel program in this manner, follow these steps:

1. Build your kernel program into the kernel on a test system.
2. Set up the `kdebug` debugger, as described in Section 2.3.
3. Issue the `dbx -remote` command on a remote build system, supplying the pathname of the kernel running on the test system.
4. Set breakpoints and enter `dbx` commands as you normally would. Section 2.1 describes some of the commands that are useful during kernel debugging. For general information about using `dbx`, see the *Programmer's Guide*.

The system also provides the `kdbx` debugger, which is designed especially for debugging kernel code. This debugger contains a number of special commands, called extensions, that allow you to display kernel data structures in a readable format. Section 2.2 describes using `kdbx` and its extensions. (You cannot use the `kdbx` debugger with the `kdebug` debugger.)

Another feature of `kdbx` is that you can customize it by writing your own extensions. The system contains a set of `kdbx` library routines that you can use to create extensions that display kernel data structures in ways that are meaningful to you. Chapter 3 describes writing `kdbx` extensions.

1.3 Debugging the Running Kernel

When you have problems with a process or set of processes, you can attempt to identify the problem by debugging the running kernel. You might also invoke the debugger on the running kernel to examine the values assigned to system parameters. (You can modify the value of the parameters using the debugger, but this practice can cause problems with the kernel and should be avoided.)

You use the `dbx` or `kdbx` debugger to examine the state of processes running on your system and to examine the value of system parameters. The `kdbx` debugger provides special commands, called extensions, that you can use to display kernel data structures. (Section 2.2.3 describes the extensions.)

To examine the state of processes, you invoke the debugger (as described in Section 2.1 or Section 2.2) using the following command:

```
# dbx -k /vmunix /dev/mem
```

This command invokes `dbx` with the kernel debugging flag, `-k`, which maps kernel addresses to make kernel debugging easier. The `/vmunix` and `/dev/mem` parameters cause the debugger to operate on the running kernel.

Once in the `dbx` environment, you use `dbx` commands to display process IDs and trace execution of processes. You can perform the same tasks using the `kdbx` debugger. The following example shows the `dbx` command you use to display process IDs:

```
(dbx) kps
      PID  COMM
00000  kernel idle
00001  init
00014  kloadsrv
00016  update
:
```

If you want to trace the execution of the `kloadsrv` daemon, use the `dbx` command to set the `$pid` symbol to the process ID of the `kloadsrv` daemon. Then, enter the `t` command:

```
(dbx) set $pid = 14
(dbx) t
> 0 thread_block() ["/usr/sde/build/src/kernel/kern/sched_prim.c":1623, 0xfffffc0000\
43d77c]
  1 mpsleep(0xffffffff92586f00, 0x11a, 0xfffffc0000279cf4, 0x0, 0x0) ["/usr/sde/build\
/src/kernel/bsd/kern_synch.c":411, 0xfffffc000040adc0]
  2 sosleep(0xffffffff92586f00, 0x1, 0xfffffc000000011a, 0x0, 0xfffffffff81274210) ["/usr/sde\
/build/src/kernel/bsd/uipc_socket2.c":654, 0xfffffc0000254ff8]
  3 sosbwait(0xffffffff92586f60, 0xffffffff92586f00, 0x0, 0xffffffff92586f00, 0x10180) ["/usr\
/sde/build/src/kernel/bsd/uipc_socket2.c":630, 0xfffffc0000254fe4]
  4 soreceive(0x0, 0xffffffff9a64f658, 0xffffffff9a64f680, 0x8000004300000000, 0x0) ["/usr/sde\
/build/src/kernel/bsd/uipc_socket.c":1297, 0xfffffc0000253338]
  5 recvrit(0xfffffc0000456fe8, 0xffffffff9a64f718, 0x14000c6d8, 0xffffffff9a64f8b8,\
0xfffffc000043d724) ["/usr/sde/build/src/kernel/bsd/uipc_syscalls.c":1002,\
0xfffffc00002574f0]
  6 recvfrom(0xfffffffff81274210, 0xffffffff9a64f8c8, 0xffffffff9a64f8b8, 0xffffffff9a64f8c8,\
0xfffffc0000457570) ["/usr/sde/build/src/kernel/bsd/uipc_syscalls.c":860,\
0xfffffc000025712c]
  7 orecvfrom(0xffffffff9a64f8b8, 0xffffffff9a64f8c8, 0xfffffc0000457570, 0x1, 0xfffffc0000456fe8)\
["/usr/sde/build/src/kernel/bsd/uipc_syscalls.c":825, 0xfffffc000025708c]
  8 syscall(0x120024078, 0xffffffffffffffff, 0xffffffffffffffff, 0x21, 0x7d) ["/usr/sde\
/build/src/kernel/arch/alpha/syscall_trap.c":515, 0xfffffc0000456fe4]
  9 _Xsyscall(0x8, 0x12001acb8, 0x14000eed0, 0x4, 0x1400109d0) ["/usr/sde/build\
/src/kernel/arch/alpha/locore.s":1046, 0xfffffc00004486e4]
(dbx) exit
```

Often, looking at the trace of a process that is hanging or has unexpectedly stopped running reveals the problem. Once you find the problem, you can modify system parameters, restart daemons, or take other corrective actions.

For more information about the commands you can use to debug the running kernel, see Section 2.1 and Section 2.2.

1.4 Analyzing a Crash Dump File

If your system crashes, you can often find the cause of the crash by using `dbx` or `kdbx` to debug or analyze a crash dump file.

The operating system can crash because one of the following occurs:

- Hardware exception
- Software panic
- Hung system

When a system hangs, it is often necessary to force the system to create dumps that you can analyze to determine why the system hung. The *System Administration* manual describes the procedure for forcing a crash dump of a hung system.

- Resource exhaustion

The system crashes or hangs because it cannot continue executing. Normally, even in the case of a hardware exception, the operating system detects the problem. (For example a machine-checking routine might discover a hardware problem and begin the process of crashing the system.) In general, the operating system performs the following steps when it detects a problem from which it cannot recover:

1. It calls the system `panic` function.

The `panic` function saves the contents of registers and sends the panic string (a message describing the reason for the system panic) to the error logger and the console terminal.

If the system is a Symmetric Multiprocessing (SMP) system, the `panic` function notifies the other CPUs in the system that a panic has occurred. The other CPUs then also execute the `panic` function and record the following panic string:

```
cpu_ip_intr: panic request
```

Once each CPU has recorded the system panic, execution continues only on the master CPU. All other CPUs in the SMP system stop execution.

2. It calls the system `boot` function.

The `boot` function records the stack.

3. It calls the `dump` function.

The `dump` function copies core memory into swap partitions and the system stops running or the reboot process begins. Console environment variables control whether the system reboots automatically. (The *System Administration* manual describes these environment variables.)

At system reboot time, the copy of core memory saved in the swap partitions is copied into a file, called a crash dump file. You can analyze the crash dump file to determine what caused the crash. By default, the crash dump is a partial (rather than full) dump and is in compressed form. For complete information about managing crash dumps and crash dump files, including how to change default settings, see the *System Administration* manual. For examples of analyzing crash dump files, see Chapter 4.

2

Kernel Debugging Utilities

The Tru64 UNIX system provides several tools you can use to debug the kernel and kernel programs. The Ladebug debugger (available as an option) is also capable of debugging the kernel.

This chapter describes three debuggers and a utility for analyzing crash dumps:

- The `dbx` debugger, which is described for kernel debugging in Section 2.1. (For general `dbx` user information, see the *Programmer's Guide*.)

You can use the `dbx` debugger to display the values of kernel variables and kernel structures. However, you must understand the structures and be prepared to follow the address links to find the information you need. You cannot use `dbx` alone to control execution of the running kernel, for example by setting breakpoints.

- The `kdbx` debugger, which is described in Section 2.2.

The `kdbx` debugger is an interface to `dbx` that is tailored specifically to debugging kernel code. The `kdbx` debugger has knowledge of the structure of kernel data and so displays kernel data in a readable format. Also, `kdbx` is extensible, allowing you to create commands that are tailored to your kernel-debugging needs. (Chapter 3 describes how to tailor the `kdbx` debugger.) However, you cannot use `dbx` command line editing features when you use the `kdbx` debugger.

- The `kdebug` debugger, which is described in Section 2.3.

The `kdebug` debugger is a kernel-debugging program that resides inside the kernel. Working with a remote version of the `dbx` debugger, the `kdebug` debugger allows you to set breakpoints in and control the execution of kernel programs and the kernel.

- The `crashdc` utility, which is described in Section 2.4.

The `crashdc` utility is a crash dump analysis tool. This utility is useful when you need to determine why the system is hanging or crashing.

The sections that follow describe how to use these tools to debug the kernel and kernel programs.

Note

Starting with Tru64 UNIX Version 5.0, all the previously mentioned tools can be used with compressed (`vmzcore.n`) and uncompressed (`vmcore.n`) crash dump files. Older versions of these tools can read only `vmcore.n` files. If you are using an older version of a tool, use the `expand_dump` utility to produce a `vmcore.n` file from a `vmzcore.n` file. For more information about compressed and uncompressed crash dump files, see `expand_dump(8)` and the *System Administration* manual.

2.1 The dbx Debugger

The `dbx` debugger is a symbolic debugger that allows you to examine, modify, and display the variables and data structures found in stripped or nonstripped kernel images.

The following sections describe how to invoke the `dbx` debugger for kernel debugging (Section 2.1.1) and how to use its commands to perform tasks such as the following:

- Debugging stripped images (Section 2.1.2)
- Specifying the location of loadable modules for crash dumps (Section 2.1.3)
- Examining memory contents (Section 2.1.4)
- Displaying the values of kernel variables, and the value and format of kernel data structures (Section 2.1.5)
- Displaying the format of a data structure (Section 2.1.6)
- Debugging multiple threads (Section 2.1.7)
- Examining the exception frame (Section 2.1.8)
- Examining the user program stack (Section 2.1.9)
- Extracting the preserved message buffer (Section 2.1.10)
- Debugging on SMP systems (Section 2.1.11)

For more information on `dbx`, see the *Programmer's Guide*.

2.1.1 Invoking the dbx Debugger for Kernel Debugging

To debug kernel code with the `dbx` debugger, you use the `-k` flag. This flag causes `dbx` to map memory addresses. When you use the `dbx -k` command, the debugger operates on two separate files that reflect the current state of the kernel that you want to examine. These files are as follows:

- The disk version of the executable kernel image

- The system core memory image

These files may be files from a running system, such as `/vmunix` and `/dev/mem`, or dump files, such as `vmunix.n` and `vmzcore.n` (compressed) or `vmcore.n` (uncompressed). By default, crash dump files are created in the `/var/adm/crash` directory (see the *System Administration* manual).

Note

You might need to be the superuser (`root` login) to examine the running system or crash dump files produced by `savecore`. Whether you need to be the superuser depends on the directory and file protections for the files you attempt to examine with the `dbx` debugger.

Use the following `dbx` command to examine the running system:

```
# dbx -k /vmunix /dev/mem
```

Use a `dbx` command similar to the following to examine a compressed or uncompressed crash dump file, respectively:

```
# dbx -k vmunix.1 vmzcore.1
# dbx -k vmunix.1 vmcore.1
```

The version number (.1, in this example) is determined by the value contained in the `bounds` file, which is located in the same directory as the dump files.

2.1.2 Debugging Stripped Images

By default, the kernel is compiled with a debugging flag that does not strip all of the symbol table information from the executable kernel image. The kernel is also partially optimized during the compilation process by default. If the kernel or any other file is fully optimized and stripped of all symbol table information during compilation, your ability to debug the file is greatly reduced. However, the `dbx` debugger provides commands to aid you in debugging stripped images.

When you attempt to display the contents of a symbol during a debugging session, you might encounter messages such as the following:

```
No local symbols.
Undefined symbol.
Inactive symbol.
```

These messages might indicate that you are debugging a stripped image.

To see the contents of all symbols during a debugging session, you can leave the debugging session, rebuild all stripped modules (but do not strip them), and reenter the debugging session. However, on certain occasions, you might

want to add a symbol table to your current debugging session rather than end the session and start a new one. To add a symbol table to your current debugging session, follow these steps:

1. Go to a window other than the one in which the debugger is running, or put the debugger in the background, and rebuild the modules for which you need a symbol table.
2. Once the modules build correctly, use the `ostrip` command to strip a symbol table out of the resulting executable file. For example, if your executable file is named `kernel_program`, issue a command such as the following one:

```
% /usr/ucb/ostrip -t kernel_program
```

The `-t` flag causes the `ostrip` command to produce two files. One, named `kernel_program`, is the stripped executable image. The other, named `kernel_program.stb`, contains the symbol table information for the `kernel_program` module. (For more information about the `ostrip` command, see `ostrip(1)`.)

3. Return to the debugging session and add the symbol table file by issuing the `dbx` command `stbadd` as follows:

```
dbx> stbadd kernel_program.stb
```

You can specify an absolute or relative pathname on the `stbadd` command line.

Once you issue this command, you can display the contents of symbols included in the symbol table just as if you had built the module you are debugging without stripping.

You can also delete symbol tables from a debugging session using the `dbx` command `stbdel`. For more information about this command, see `dbx(1)`.

2.1.3 Specifying the Location of Loadable Modules for Crash Dumps

When a crash dump occurs, the location of any loadable modules used by the kernel is recorded in the crash dump file, enabling `dbx` to find the modules. If the version of a loadable module that was running when the crash occurred is moved to a different location, `dbx` will not find it. You can specify the directory path where `dbx` should look for loadable modules by using any one of the following methods (see `dbx(1)` for complete details):

- On the `dbx` command line, specify the directory path with the `-module_path` option. For example:

```
# dbx -k vmunix.1 vmzcore.1 -module_path /project4/mod_dir
```

- Before invoking `dbx`, set the environment variable `DBX_MODULE_PATH`. For example:

```
# setenv DBX_MODULE_PATH /project4/mod_dir
```

- During the dbx session, if you want to load a module dynamically, first set the `$module_path` dbx variable and then use the `addobj` command to load the module, as in the following example:

```
(dbx) set $module_path /project4/mod_dir
(dbx) addobj kmodC
```

To verify that modules are being loaded from the correct location, turn on verbose module-loading using any one of the following methods:

- Specify the `-module_verbose` dbx command option.
- Set the `DBX_MODULE_VERBOSE` environment variable to any integer value.
- Set the `$module_verbose` dbx variable to a nonzero value.

2.1.4 Examining Memory Contents

To examine memory contents with dbx, use the following syntax:

```
address/count[mode]
```

The *count* argument specifies the number of items that the debugger displays at the specified *address*, and the *mode* argument determines how dbx displays memory. If you omit the *mode* argument, the debugger uses the previous mode. The initial default mode is `x` (hexadecimal). Table 2–1 lists the dbx address modes.

Table 2–1: The dbx Address Modes

Mode	Description
b	Displays a byte in octal.
c	Displays a byte as a character.
d	Displays a short word in decimal.
D	Displays a long word in decimal.
f	Displays a single precision real number.
g	Displays a double precision real number.
i	Displays machine instructions.
n	Displays data in typed format.
o	Displays a short word in octal.
O	Displays a long word in octal.
s	Displays a string of characters that ends in a null.
x	Displays a short word in hexadecimal.
X	Displays a long word in hexadecimal.

The following examples show how to use dbx to examine kernel images:

```
(dbx) _realstart/X
fffffc00002a4008: c020000243c4153e
(dbx) _realstart/i
[_realstart:153, 0xfffffc00002a4008] subq    sp, 0x20, sp
(dbx) _realstart/10i
[_realstart:153, 0xfffffc00002a4008] subq    sp, 0x20, sp
[_realstart:154, 0xfffffc00002a400c] br     r1, 0xfffffc00002a4018
[_realstart:156, 0xfffffc00002a4010] call_pal    0x4994e0
[_realstart:157, 0xfffffc00002a4014] bgt    r31, 0xfffffc00002a3018
[_realstart:171, 0xfffffc00002a4018] ldq    gp, 0(r1)
[_realstart:172, 0xfffffc00002a401c] stq    r31, 24(sp)
[_realstart:177, 0xfffffc00002a4020] bis    r16, r31, r9
[_realstart:178, 0xfffffc00002a4024] bis    r17, r31, r10
[_realstart:179, 0xfffffc00002a4028] bis    r18, r31, r11
[_realstart:181, 0xfffffc00002a402c] bis    r19, r31, r12
```

2.1.5 Printing the Values of Variables and Data Structures

You can use the `print` command to examine values of variables and data structures. The `print` command has the following syntax:

print *expression*

p *expression*

For example:

```
(dbx) print utsname
struct {
    sysname = "OSF1"
    nodename = "system.dec.com"
    release = "V5.0"
    version = "688.2"
    machine = "alpha"
}
```

Note that dbx has a default alias of `p` for `print`:

```
(dbx) p utsname
```

2.1.6 Displaying a Data Structure Format

You can use the `whatis` command to display the format for many of the kernel data structures. The `whatis` command has the following syntax:

whatis *type name*

The following example displays the `itimerval` data structure:

```
(dbx) whatis struct itimerval
struct itimerval {
    struct timeval {
        int tv_sec;
        int tv_usec;
    } it_interval;
```

```

    struct timeval {
        int tv_sec;
        int tv_usec;
    } it_value;
};

```

2.1.7 Debugging Multiple Threads

You can use the `dbx` debugger to examine the state of the kernel's threads with the querying and scoping commands described in this section. You use these commands to show process and thread lists and to change the debugger's context (by setting its current process and thread variables) so that a stack trace for a particular thread can be displayed. Use these commands to examine the state of the kernel's threads:

<code>print \$tid</code>	Display the thread ID of the current thread
<code>print \$pid</code>	Display the process ID of the current process
<code>where</code>	Display a stack trace for the current thread
<code>tlist</code>	Display a list of kernel threads for the current process
<code>kps</code>	Display a list of processes (not available when used with <code>kdebug</code>)
<code>set \$pid=<i>process_id</i></code>	Change the context to another process (a process ID of 0 changes context to the kernel)
<code>tset <i>thread_id</i></code>	Change the context to another thread
<code>tstack</code>	Displays the stack trace for all threads.

2.1.8 Examining the Exception Frame

When you work with a crash dump file to debug your code, you can use `dbx` to examine the exception frame. The exception frame is a stack frame created during an exception. It contains the registers that define the state of the routine that was running at the time of the exception. Refer to the

/usr/include/machine/reg.h header file to determine where registers are stored in the exception frame.

The `savedefp` variable contains the location of the exception frame. (Note that no exception frames are created when you force a system to dump, as described in the *System Administration* manual.) The following example shows an example exception frame:

```
(dbx) print savedefp/33X
ffffffff9618d940: 0000000000000000 ffffffc000046f888
ffffffff9618d950: ffffffff86329ed0 0000000079cd612f
ffffffff9618d960: 000000000000007d 0000000000000001
ffffffff9618d970: 0000000000000000 ffffffc000046f4e0
ffffffff9618d980: 0000000000000000 ffffffff9618a2f8
ffffffff9618d990: 0000000140012b20 0000000000000000
ffffffff9618d9a0: 000000014002ee10 0000000000000000
ffffffff9618d9b0: 00000001400075e8 0000000140026240
ffffffff9618d9c0: ffffffff9618daf0 ffffffff8635af20
ffffffff9618d9d0: ffffffff9618dac0 00000000000001b0
ffffffff9618d9e0: fffffc00004941b8 0000000000000000
ffffffff9618d9f0: 0000000000000001 fffffc000028951c
ffffffff9618da00: 0000000000000000 00000000000000ff
ffffffff9618da10: 0000000140026240 0000000000000000
ffffffff9618da20: 0000000000000000 fffffc000047acd0
ffffffff9618da30: 0000000000901402 000000000001001
ffffffff9618da40: 0000000000002000
```

2.1.9 Examining the User Program Stack

When debugging a crash dump with `dbx`, you can examine the call stack of the user program whose execution precipitated the kernel crash. To examine a crash dump and also view the user program stack, you must invoke `dbx` using the following command syntax:

```
dbx -k vmunix.n vm[z]core.n path/user-program
```

The version number (*n*) is determined by the value contained in the `bounds` file, which is located in the same directory as the dump files. The `user-program` parameter specifies the user program executable.

The crash dump file must contain a full crash dump. For information on setting system defaults for full or partial crash dumps, see the *System Administration* manual. You can use the `assign` command in `dbx`, as shown in the following example, to temporarily specify a full crash dump. This setting stays in effect until the system is rebooted.

```
# dbx -k vmunix.3
dbx version 5.0
.
.
.
(dbx) assign partial_dump=0
```


To specify a full crash dump permanently so that this setting remains in effect after a reboot, use the patch command in dbx, as shown in the following example:

```
(dbx) patch partial_dump=0
```

With either command, a `partial_dump` value of 1 specifies a partial dump.

The following example shows how to examine the state of a user program named `test1` that purposely precipitated a kernel crash with a `syscall` after several recursive calls:

```
# dbx -k vmunix.1 vmzcore.1 /usr/proj7/test1
dbx version 5.0
Type 'help' for help.

stopped at [boot:1890 ,0xfffffc000041ebe8]      Source not available

warning: Files compiled -g3: parameter values probably wrong
(dbx) where 1
> 0 boot() ["/../../../../src/kernel/arch/alpha/machdep.c":1890,
0xfffffc000041ebe8]
  1 panic(0xfffffc000051e1e0, 0x8, 0x0, 0x0, 0xffffffff888c3a38)
["../../../../src/kernel/bsd/subr_prf.c":824, 0xfffffc0000281974]
  2 syscall(0x2d, 0x1, 0xffffffff888c3ce0, 0x9aa1e0000000, 0x0)
["../../../../src/kernel/arch/alpha/syscall_trap.c":593, 0xfffffc0000423be4]
  3 _Xsyscall(0x8, 0x3ff8010f9f8, 0x140008130, 0xaa, 0x3ffc0097b70)
["../../../../src/kernel/arch/alpha/locore.s":1409, 0xfffffc000041b0f4]
  4 __syscall(0x0, 0x0, 0x0, 0x0, 0x0) [0x3ff8010f9f4]
  5 justtryme(scall = 170, cpu = 0, levels = 25) ["test1.c":14,
0x120001310]
  6 recurse(inbox = (...)) ["test1.c":28, 0x1200013c4]
  7 recurse(inbox = (...)) ["test1.c":30, 0x120001400]
  8 recurse(inbox = (...)) ["test1.c":30, 0x120001400]
  9 recurse(inbox = (...)) ["test1.c":30, 0x120001400]
  .
  .
  .
 30 recurse(inbox = (...)) ["test1.c":30, 0x120001400]
 31 main(argc = 3, argv = 0x11ffffd08) ["test1.c":52, 0x120001518]
(dbx) up 8 2
recurse: 30 if (r.a[2] > 0) recurse(r);
(dbx) print r 3
struct {
  a = {
    [0] 170
    [1] 0
    [2] 2
    [3] 0
    .
    .
    .
  }
}
(dbx) print r.a[511] 4
25
(dbx)
```

- 1 The `where` command displays the kernel stack followed by the user program stack at the time of the crash. In this case, the kernel stack has 4 activation levels; the user program stack starts with the fifth level and includes several recursive calls.

- ❷ The `up 8` command moves the debugging context 8 activation levels up the stack to one of the recursive calls within the user program code.
- ❸ The `print r` command displays the current value of the variable `r`, which is a structure of array elements. Full symbolization is available for the user program, assuming it was compiled with the `-g` option.
- ❹ The `print r.a[511]` command displays the current value of array element 511 of structure `r`.

2.1.10 Extracting the Preserved Message Buffer

The preserved message buffer (`pmsgbuf`) contains information such as the firmware version, operating system version, `pc` value, and device configuration. You can use `dbx` to extract the preserved message buffer from a running system or dump files. For example:

```
(dbx) print *pmsgbuf
struct {
    msg_magic = 405601
    msg_bufx = 1537
    msg_bufr = 1537
    msg_bufc = "Alpha boot: available memory from 0x7c6000 to 0x6000000
Tru64 UNIX V5.0; Sun Jan 03 11:20:36 EST 1999
physical memory = 96.00 megabytes.
available memory = 84.57 megabytes.
using 360 buffers containing 2.81 megabytes of memory
tc0 at nexus
scc0 at tc0 slot 7
asc0 at tc0 slot 6
rz1 at scsi0 target 1 lun 0 (LID=0) (DEC      RZ25      (C) DEC 0700)
rz2 at scsi0 target 2 lun 0 (LID=1) (DEC      RZ25      (C) DEC 0700)
rz3 at scsi0 target 3 lun 0 (LID=2) (DEC      RZ26      (C) DEC T384)
rz4 at scsi0 target 4 lun 0 (LID=3) (DEC      RRD42      (C) DEC 4.5d)
tz5 at scsi0 target 5 lun 0 (DEC      TLZ06      (C) DEC 0374)
scsil at tc0 slot 7
fb0 at tc0 slot 8
  1280X1024
ln0: DEC LANCE Module Name: PMAD-BA
ln0 at tc0 slot 7
:
:
```

2.1.11 Debugging on SMP Systems

Debugging in an SMP environment can be difficult because an SMP system optimized for performance keeps the minimum of lock debug information.

The Tru64 UNIX system supports a lock mode to facilitate debugging SMP locking problems. The lock mode is implemented in the `lockmode` boot time system attribute. By default, the `lockmode` attribute is set to a value between 0 and 3, depending upon whether the system is an SMP system and whether the `RT_PREEMPTION_OPT` attribute is set. (This attribute optimizes system performance.)

For debugging purposes, set the `lockmode` attribute to 4. Follow these steps to set the `lockmode` attribute to 4:

1. Create a stanza-formatted file named, for example, `generic.stanza` that appears as follows:

```
generic:
    lockmode=4
```

The contents of this file indicate that you are modifying the `lockmode` attribute of the `generic` subsystem.

2. Add the new definition of `lockmode` to the `/etc/sysconfigtab` database:

```
# sysconfigdb -a -f generic.stanza generic
```

3. Reboot your system.

Some of the debugging features provided with `lockmode` set to 4 are as follows:

- Automatic lock hierarchy checking and minimum `spl` checking when any kernel lock is acquired (assuming a `lockinfo` structure exists for the lock class in question). This checking helps you find potential deadlock situations.
- Lock initialization checking.
- Additional debug information maintenance, including information about simple and complex locks.

For simple locks, the system records an array of the last 32 simple locks which were acquired on the system (`slock_debug`). The system creates a `slock_debug` array for each CPU in the system.

For complex locks, the system records the locks owned by each thread in the thread structure (up to eight complex locks).

To get a list of the complex locks a thread is holding use these commands:

```
# dbx -k /vmunix
(dbx) print thread->lock_addr
{
    [0] 0xe4000002a67e0030
    [1] 0xc3e0005b47ff0411
    [2] 0xb67e0030a6130048
    [3] 0xa67e0030d34254e5
    [4] 0x279f0200481e1617
    [5] 0x4ae33738a7730040
    [6] 0x477c0101471c0019
    [7] 0xb453004047210402
}
(dbx) print slock_debug
{
```

```

    [0] 0xfffffc000065c580
    [1] 0xfffffc000065c780
}

```

- Lock statistics are recorded to allow you to determine what kind of contention you have on a particular lock. Use the `kdbx lockstats` extension as shown in the following example to display lock statistics:

```

# kdbx /vmunix
(kdbx) lockstats
Lockstats      li_name          cpu count  tries    misses %misses waitsum    waitmax waitmin  trmax
=====
k0x00657d40    inode.i_io_lock  1  1784    74268   1936 2.61    110533    500      6      10
k0x00653400    nfs_daemon_lock  0   1         7         0 0.00      0         0         0         0
k0x00657d80    nfs_daemon_lock  1   1         0         0 0.00      0         0         0         0
k0x00653440    lk_lmf           0   1         0         0 0.00      0         0         0         0
k0x00657dc0    lk_lmf           1   1         2         0 0.00      0         0         0         0
k0x00653480    procfs_global_lock 0   1         3         0 0.00      0         0         0         0
k0x00657e00    procfs_global_lock 1   1         5         0 0.00      0         0         0         0
k0x006534c0    procfs_pr_trace_lock 0  40         0         0 0.00      0         0         0         0
k0x00657e40    procfs_pr_trace_lock 1  40         0         0 0.00      0         0         0         0

```

2.2 The kdbx Debugger

The `kdbx` debugger is a crash analysis and kernel debugging tool; it serves as a front end to the `dbx` debugger. The `kdbx` debugger is extensible, customizable, and insensitive to changes to offsets and field sizes in structures. The only dependencies on kernel header files are for bit definitions in flag fields.

The `kdbx` debugger has facilities for interpreting various symbols and kernel data structures. It can format and display these symbols and data structures in the following ways:

- In a predefined form as specified in the source code modules that currently accompany the `kdbx` debugger
- As defined in user-written source code modules according to a standardized format for the contents of the `kdbx` modules

All `dbx` commands (except signals such as `Ctrl/P`) are available when you use the `kdbx` debugger. In general, `kdbx` assumes hexadecimal addresses for commands that perform input and output.

As with `dbx`, you can use `kdbx` to examine the call stack of the user program whose execution precipitated a kernel crash (see Section 2.1.9).

The sections that follow explain using `kdbx` to debug kernel programs.

2.2.1 Beginning a kdbx Session

Using the `kdbx` debugger, you can examine the running kernel or dump files created by the `savecore` utility. In either case, you examine an object file and a core file. For running systems, these files are usually `/vmunix` and

/dev/mem, respectively. By default, crash dump files are created in the /var/adm/crash directory (see the *System Administration* manual).

Use the following `kdbx` command to examine a running system:

```
# kdbx -k /vmunix /dev/mem
```

Use a `kdbx` command similar to the following to examine a compressed or uncompressed crash dump file, respectively:

```
# kdbx -k vmunix.1 vmzcore.1
# kdbx -k vmunix.1 vmcore.1
```

The version number (.1 in this example) is determined by the value contained in the `bounds` file, which is located in the same directory as the dump files.

To examine a crash dump file and also view the call stack of the user program whose execution precipitated the kernel crash, you must invoke `kdbx` using the following command syntax:

```
kdbx -k vmunix.n vm[z]core.n path/user-program
```

For more information, see Section 2.1.9.

When you begin a debugging session, `kdbx` reads and executes the commands in the system initialization file `/var/kdbx/system.kdbxrc`. The initialization file contains setup commands and alias definitions. (For a list of `kdbx` aliases, see the `kdbx(1)` reference page.) You can further customize the `kdbx` environment by adding commands and aliases to:

- The `/var/kdbx/site.kdbxrc` file
This file contains customized commands and alias definitions for a particular system.
- The `~/.kdbxrc` file
This file contains customized commands and alias definitions for a specific user.
- The `./kdbxrc` file
This file contains customized commands and alias definitions for a specific project. This file must reside in the current working directory when `kdbx` is invoked.

2.2.2 The `kdbx` Debugger Commands

The `kdbx` debugger provides the following commands:

```
alias [name] [command-string]
```

Sets or displays aliases. If you omit all arguments, `alias` displays all aliases. If you specify the variable `name`, `alias` displays the alias for `name`, if one exists. If you specify `name` and `command-string`, `alias` establishes `name` as an alias for `command-string`.

`context proc | user`

Sets context to the user's aliases or the extension's aliases. This command is used only by the extensions.

`coredata start_address end_address`

Dumps, in hexadecimal, the contents of the core file starting at *start_address* and ending before *end_address*.

`dbx command-string`

Passes the *command-string* to dbx. Specifying dbx is optional; if kdbx does not recognize a command, it automatically passes that command to dbx. See the dbx(1) reference page for a complete description of dbx commands.

`help [-long] [args]`

Prints help text.

`pr [flags] [extensions] [arguments]`

Executes an extension and gives it control of the kdbx session until it quits. You specify the name of the extension in *extension* and pass arguments to it in *arguments*.

- | | |
|---|--|
| <code>-debug</code> | Causes kdbx to display input to and output from the extension on the screen. |
| <code>-pipe in_pipe
out_pipe</code> | Used in conjunction with the dbx debugger for debugging extensions. See Chapter 3 for information on using the <code>-pipe</code> flag. |
| <code>-print_output</code> | Causes the output of the extension to be sent to the invoker of the extension without interpretation as kdbx commands. |
| <code>-redirect_output</code> | Used by extensions that execute other extensions to redirect the output from the called extensions; otherwise, the user receives the output. |
| <code>-tty</code> | Causes kdbx to communicate with the subprocess through a terminal line instead of pipes. If you specify the <code>-pipe</code> flag, <code>proc</code> ignores it. |

```
print string
```

Displays *string* on the terminal. If this command is used by an extension, the terminal receives no output.

```
quit
```

Exits the kdbx debugger.

```
source [-x] [file(s)]
```

Reads and interprets files as kdbx commands in the context of the current aliases. If you specify the `-x` flag, the debugger displays commands as they are executed.

```
unalias name
```

Removes the alias, if any, from *name*.

The kdbx debugger contains many predefined aliases, which are defined in the kdbx startup file `/var/kdbx/system.kdbxrc`.

2.2.3 Using kdbx Debugger Extensions

In addition to its commands, the kdbx debugger provides extensions. You execute extensions using the kdbx command `pr`. For example, to execute the `arp` extension, you enter this command:

```
kdbx> pr arp
```

Some extensions are provided with your Tru64 UNIX system and reside in the `/var/kdbx` directory. Aliases for each of these extensions are also provided that let you omit the `pr` command from an extension command line. Thus, another way to execute the `arp` extension is to enter the following command:

```
kdbx> arp
```

This command has the same effect as the `pr arp` command.

You can create your own kdbx extensions as described in Chapter 3.

For extensions that display addresses as part of their output, some use a shorthand notation for the upper 32-bits of an address to keep the output readable. The following table lists the notation for each address type.

Notation	Address Type	Replaces	Example
v	virtual	ffffffff	v0x902416f0
e	virtual	fffffffe	e0x12340000

Notation	Address Type	Replaces	Example
k	kseg	ffffffc00	k0x00487c48
u	user space	00000000	u0x86406200
?	Unrecognized or random type		?0x3782cc33

The sections that follow describe the kdbx extensions that are supplied with your system.

2.2.3.1 Displaying the Address Resolution Protocol Table

The arp extension displays the contents of the address resolution protocol (arp) table. The arp extension has the following form:

arp [-]

If you specify the optional hyphen (-), arp displays the entire arp table; otherwise, it displays those entries that have nonzero values in the iaddr.s_addr and at_flags fields.

For example:

```
(kdbx) arp
      NAME          BUCK SLOT      IPADDR          ETHERADDR  MHOLD  TIMER  FLAGS
=====
sys1.zk3.dec.com  11    0  16.140.128.4  170.0.4.0.91.8  0    450    3
sys2.zk3.dec.com  18    0  16.140.128.1  0.0.c.1.8.e8    0    194    3
sys3.zk3.dec.com  31    0  16.140.128.6  8.0.2b.24.23.64  0    539   103
```

2.2.3.2 Performing Commands on Array Elements

The array_action extension performs a command action on each element of an array. This extension allows you to step through any array in the operating system kernel and display specific components or values as described in the list of command flags.

This extension has the following format:

array_action "type" length start_address [flags] command

The arguments to the array_action extension are as follows:

"type" The type of address of an element in the specified array.

length The number of elements in the specified array.

start_address The address of an array. The address can be specified as a variable name or a number. The more common syntax or notation used to refer

to the *start_address* is usually of the form `&arrayname[0]`.

flags If the you specify the `-head` flag, the next argument appears as the table header.

If the you specify the `-size` flag, the next argument is used as the array element size; otherwise, the size is calculated from the element type.

If the you specify the `-cond` flag, the next argument is used as a filter. It is evaluated by `dbx` for each array element, and if it evaluates to `TRUE`, the action is taken on the element. The same substitutions that are applied to the command are applied to the condition.

command The `kdbx` or `dbx` command to perform on each element of the specified array.

Note

The `kdbx` debugger includes several aliases, such as `file_action`, that may be easier to use than using the `array_action` extension directly.

Substitutions similar to `printf` can be performed on the command for each array element. The possible substitutions are as follows:

Conversion Character	Description
<code>%a</code>	Address of element
<code>%c</code>	Cast of address to pointer to array element
<code>%i</code>	Index of element within the array
<code>%s</code>	Size of element
<code>%t</code>	Type of pointer to element

For example:

```
(kdbx) array_action "struct kernargs *" 11 &kernargs[0] p %c.name
0xffffffffc00004737f8 = "askme"
0xffffffffc0000473800 = "bufpages"
0xffffffffc0000473810 = "nbuf"
0xffffffffc0000473818 = "memlimit"
0xffffffffc0000473828 = "pmap_debug"
```

```

0xffffffff0000473838 = "syscalltrace"
0xffffffff0000473848 = "boothowto"
0xffffffff0000473858 = "do_virtual_tables"
0xffffffff0000473870 = "netblk"
0xffffffff0000473878 = "zalloc_physical"
0xffffffff0000473888 = "trap_debug"
(kdbx)

```

2.2.3.3 Displaying the Buffer Table

The `buf` extension displays the buffer table. This extension has the following format:

```
buf [addresses -free | -all]
```

If you omit arguments, the debugger displays the buffers on the hash list.

If you specify addresses, the debugger displays the buffers at those addresses. Use the `-free` flag to display buffers on the free list. Use the `-all` flag to display first buffers on the hash list, followed by buffers on the free list.

For example:

```

(kdbx) buf
BUF MAJ MIN BLOCK COUNT SIZE RESID VNO FWD BACK FLAGS
=====
Bufs on hash lists:
v0x904e1b30 8 2 54016 8192 8192 0 v0x902220d0 v0x904f23a8 v0x904e1d20 write cache
v0x904e21f8 8 1025 131722 1024 8192 0 v0x90279800 v0x904e3748 v0x904e22f0 write cache
v0x904e46c8 8 1025 107952 2048 8192 0 v0x90220fa8 v0x904e22f0 v0x904e23e8 read cache
v0x904e9ef0 8 2050 199216 8192 8192 0 v0x90221560 v0x904f2b68 v0x904e66c0 read cache
v0x904df758 8 1025 107968 8192 8192 0 v0x90220fa8 v0x904eac80 v0x904df378 write cache
v0x904eb538 8 2050 223840 8192 8192 0 v0x90221560 v0x904ec990 v0x904eb440 read
v0x904e5930 8 2050 379600 8192 8192 0 v0x90221560 v0x904f3fc0 v0x904ec5b0 read cache
v0x904eae70 8 2050 625392 2048 8192 0 v0x90221560 v0x904df378 v0x904e08c8 write cache
v0x904f3ec8 8 1025 18048 8192 8192 0 v0x90220fa8 v0x904dff18 v0x904e1560 write cache
:
:

```

```
(kdbx)
```

2.2.3.4 Displaying the Callout Table and Absolute Callout Table

The `callout` extension displays the callout table. This extension has the following format:

```
callout
```

For example:

```

(kdbx) callout
Processor: 0
Current time (in ticks): 615421360

FUNCTION ARGUMENT TICKS(delta)
=====
realitexpire k0x008ab220 30772
wakeuper k0x005d98e0 36541

```

wakeup	k0x0187a220	374923
thread_timeout	k0x010ee950	376286
thread_timeout	k0x0132f220	40724481
realitexpire	k0x01069950	80436086
thread_timeout	k0x01bba950	82582849

The `abscallout` extension displays the absolute callout table. This table contains callout entries with the absolute time in fractions of seconds. This extension has the following format:

abscallout

For example:

```
(kdbx) abscallout
Processor:          0
```

FUNCTION	ARGUMENT	SECONDS
psx4_tod_expire	k0x01580808	86386.734375
psx4_tod_expire	k0x01580840	172786.734375
psx4_tod_expire	k0x01580878	259186.734375
psx4_tod_expire	k0x015808b0	345586.718750
psx4_tod_expire	k0x015808e8	431986.718750
psx4_tod_expire	k0x01580920	518386.718750
psx4_tod_expire	k0x01580958	604786.750000
psx4_tod_expire	k0x01580990	691186.750000
psx4_tod_expire	k0x015809c8	777586.750000
psx4_tod_expire	k0x01580a00	863986.750000

2.2.3.5 Casting Information Stored in a Specific Address

The `cast` extension forces `dbx` to display part of memory as the specified type and is equivalent to the following command:

```
dbx print *((type) address)
```

The `cast` extension has the following format:

cast *address type*

For example:

```
(kdbx) cast 0xffffffff903e3828 char
'^@'
```

2.2.3.6 Displaying Machine Configuration

The `config` extension displays the configuration of the machine. This extension has the following format:

config

For example:

```
(kdbx) config
Bus #0 (0xfffffc000048c6a0): Name - "tc" Connected to - "nexus"
      Config 1 - tcconfl1   Config 2 - tcconfl2
      Controller "sc" (0xfffffc000048c970)
(kdbx)
```

2.2.3.7 Converting the Base of Numbers

The `convert` extension converts numbers from one base to another. This extension has the following format:

```
convert [-in[8 | 10 | 16]] [-out[2 | 8 | 10 | 16]] [args]
```

The `-in` and `-out` flags specify the input and output bases, respectively. If you omit `-in`, the input base is inferred from the arguments. The arguments can be numbers or variables.

For example:

```
(kdbx) convert -in 16 -out 10 864c2a14
2253138452
(kdbx)
```

2.2.3.8 Displaying CPU Use Statistics

The `cpustat` extension displays statistics about CPU use. Statistics displayed include percentages of time the CPU spends in the following states:

- Running user level code
- Running system level code
- Running at a priority set with the `nice()` function
- Idle
- Waiting (idle with input or output pending)

This extension has the following format:

```
cpustat [-update n] [-cpu n]
```

The `-update` flag specifies that `kdbx` update the output every *n* seconds.

The `-cpu` flag controls the CPU for which `kdbx` displays statistics. By default, `kdbx` displays statistics for all CPUs in the system.

For example:

```
(kdbx) cpustat
Cpu  User (%)   Nice (%) System (%)  Idle (%)  Wait (%)
=====
  0    0.23     0.00    0.08    99.64    0.05
  1    0.21     0.00    0.06    99.68    0.05
```

2.2.3.9 Disassembling Instructions

The `dis` extension disassembles some number of instructions. This extension has the following format:

dis *start-address* [*num-instructions*]

The *num-instructions*, argument specifies the number of instructions to be disassembled. The *start-address* argument specifies the starting address of the instructions. If you omit the *num-instructions* argument, 1 is assumed.

For example:

```
(kdbx) dis 0xffffffff864c2a08 5
[., 0xffffffff864c2a08]      call_pal      0x20001
[., 0xffffffff864c2a0c]      call_pal      0x800000
[., 0xffffffff864c2a10]      ldg          $f18, -13304(r3)
[., 0xffffffff864c2a14]      bgt         r31, 0xffffffff864c2a14
[., 0xffffffff864c2a18]      call_pal      0x4573d0
(kdbx)
```

2.2.3.10 Displaying Remote Exported Entries

The `export` extension displays the exported entries that are mounted remotely. This extension has the following format:

export

For example:

```
(kdbx) export
ADDR EXPORT      MAJ  MIN  INUM      GEN  MAP  FLAGS  PATH
-----
0xffffffff863bfe40 8  4098    2  1308854383 -2   0  /cdrom
0xffffffff863bfdc0 8  2050  67619  736519799 -2   0  /usr/users/user2
0xffffffff863bfe00 8  2050  15263  731712009 -2   0  /usr/staff/user
0xffffffff863bfe80 8  1024   6528  731270099 -2   0  /mnt
```

2.2.3.11 Displaying the File Table

The `file` extension displays the file table. This extension has the following format:

file [*addresses*]

If you omit the arguments, the extension displays file entries with nonzero reference counts; otherwise, it displays the file entries located at the specified addresses.

For example:

```
(kdbx) file
Addr      Type  Ref  Msg  Fileops      f_data      Cred Offset Flags
```

```

=====
v0x90406000 file 4 0 vnops v0x90259550 v0x863d5540 68 r w
v0x90406058 file 1 0 vnops v0x9025b5b8 v0x863d5e00 4096 r
v0x904060b0 file 1 0 vnops v0x90233908 v0x863d5d60 0 r
v0x90406108 file 2 0 vnops v0x90233908 v0x863d5d60 602 w
v0x90406160 file 2 0 vnops v0x90228d78 v0x863d5b80 904 r
v0x904061b8 sock 2 0 sockops v0x863b5c08 v0x863d5c20 0 r w
v0x90406210 file 1 0 vnops v0x90239e10 v0x863d5c20 2038 r
v0x90406268 file 1 0 vnops v0x90245140 v0x863d5c20 301 w a
v0x904062c0 file 3 0 vnops v0x90227880 v0x863d5900 23 r w
v0x90406318 file 2 0 vnops v0x90228b90 v0x863d5c20 856 r
v0x90406370 sock 2 0 sockops v0x863b5a08 v0x863d5c20 0 r w
:
:

```

2.2.3.12 Displaying the udb and tcb Tables

The `inpcb` extension displays the `udb` and `tcb` tables. This extension has the following format:

inpcb [-udp] [-tcp] [*addresses*]

If you omit the arguments, `kdbx` displays both tables. If you specify the `-udp` flag or the `-tcp` flag, the debugger displays the corresponding table.

If you specify the *address* argument, the `inpcb` extension ignores the `-udp` and `-tcp` flags and displays entries located at the specified address.

For example:

```

(kdbx) inpcb -tcp
TCP:
Foreign Host  FPort      Local Host  LPort      Socket      PCB      Options
0.0.0.0      0 0.0.0.0    47621      u0x00000000 u0x00000000
system.dec.com 6000 comput.dec.com 1451 v0x8643f408 v0x863da408
system.dec.com 998 comput.dec.com 1020 v0x8643fc08 v0x863da208
system.dec.com 999 comput.dec.com 514 v0x8643ac08 v0x8643d008
system.dec.com 6000 comput.dec.com 1450 v0x863fba08 v0x863dad08
system.dec.com 1008 comput.dec.com 1021 v0x86431e08 v0x86414708
system.dec.com 1009 comput.dec.com 514 v0x86412808 v0x8643ce08
system.dec.com 6000 comput.dec.com 1449 v0x86436608 v0x86415e08
system.dec.com 6000 comput.dec.com 1448 v0x86431808 v0x863daa08
:
:
0.0.0.0      0 0.0.0.0    806 v0x863e3e08 v0x863dbe08
0.0.0.0      0 0.0.0.0    793 v0x863d1808 v0x8635a708
0.0.0.0      0 0.0.0.0    0 v0x86394408 v0x8635b008
0.0.0.0      0 0.0.0.0    1024 v0x86394208 v0x8635b108
0.0.0.0      0 0.0.0.0    111 v0x863d1e08 v0x8635b208

```

2.2.3.13 Performing Commands on Lists

The `list_action` extension performs some command on each element of a linked list. This extension provides the capability to step through any linked list in the operating system kernel and display particular components. This extension has the following format:

list_action *"type" next-field end-addr start-addr [flags] command*

The arguments to the `list_action` extension are as follows:

<i>"type"</i>	The type of an element in the specified list.
<i>next-field</i>	The name of the field that points to the next element.
<i>end-addr</i>	The value of the next field that terminates the list. If the list is NULL-terminated, the value of the <i>end-addr</i> argument is zero (0). If the list is circular, the value of the <i>end-addr</i> argument is equal to the <i>start-addr</i> argument.
<i>start_addr</i>	The address of the list. This argument can be a variable name or a number address.
<i>flags</i>	Use the <code>-head header</code> flag to display the <i>header</i> argument as the table header. Use the <code>-cond arg</code> flag to filter input as specified by <i>arg</i> . The debugger evaluates the condition for each array element, and if it evaluates to true, the action is taken on the element. The same substitutions that are applied to the command are applied to the condition.
<i>command</i>	The debugger command to perform on each element of the list.

The `kdbx` debugger includes several aliases, such as `procaddr`, that might be easier than using the `list_action` extension directly.

The `kdbx` debugger applies substitutions in the same style as `printf` substitutions for each command element. The possible substitutions are as follows:

Conversion Character	Description
<code>%a</code>	Address of an element
<code>%c</code>	Cast of an address to a pointer to a list element
<code>%i</code>	Index of an element within the list
<code>%n</code>	Name of the next field
<code>%t</code>	Type of pointer to an element

For example:

```
(kdbx) list_action "struct proc *" p_nxt 0 allproc p \  
%c.task.u_address.uu_comm %c.p_pid \  
"list_action" 1382 \  
"dbx" 1380 \  
"kdbx" 1379 \  
"dbx" 1301 \  
"kdbx" 1300 \  
"sh" 1296 \  
"ksh" 1294 \  
"csh" 1288 \  
"rlogind" 1287 \  
:  
:
```

2.2.3.14 Displaying the lockstats Structures

The `lockstats` extension displays the lock statistics contained in the `lockstats` structures. Statistics are kept for each lock class on each CPU in the system. These structures provide the following information:

- The address of the structure
- The class of lock for which lock statistics are being recorded
- The CPU for which the lock statistics are being recorded
- The number of instances of the lock
- The number of times processes have tried to get the lock
- The number of times processes have tried to get the lock and missed
- The percentage of time processes miss the lock
- The total time processes have spent waiting for the lock
- The maximum amount of time a single process has waited for the lock
- The minimum amount of time a single process has waited for the lock

The lock statistics recorded in the `lockstats` structures are dynamic.

This extension is available only when the `lockmode` system attribute is set to 4.

This extension has the following format:

```
lockstats -class name | -cpu number | -read | -sum | -total | -update n
```

If you omit all flags, `lockstats` displays statistics for all lock classes on all CPUs. The following describes the flags you can use:

<code>-class name</code>	Displays the <code>lockstats</code> structures for the specified lock class. (Use the <code>lockinfo</code> command to display information about the names of lock classes.)
<code>-cpu number</code>	Displays the <code>lockstats</code> structures for the specified CPU.
<code>-read</code>	Displays the reads, sleeps attributes, and waitsums or misses.
<code>-sum</code>	Displays summary data for all CPUs and all lock types.
<code>-total</code>	Displays summary data for all CPUs.
<code>-update n</code>	Updates the display every <i>n</i> seconds.

For example:

```
(kdbx) lockstats
Lockstats  li_name          cpu count  tries    misses %misses  waitsum  waitmax waitmin  trmax
-----
k0x00657d40  inode.i_io_lock  1  1784    74268   1936 2.61    110533   500      6      10
k0x00653400  nfs_daemon_lock  0    1        7      0 0.00     0        0      0      0
k0x00657d80  nfs_daemon_lock  1    1        0      0 0.00     0        0      0      0
k0x00653440          lk_lmf          0    1        0      0 0.00     0        0      0      0
k0x00657dc0          lk_lmf          1    1        2      0 0.00     0        0      0      0
k0x00653480  procfs_global_lock  0    1        3      0 0.00     0        0      0      0
k0x00657e00  procfs_global_lock  1    1        5      0 0.00     0        0      0      0
k0x006534c0  procfs_pr_trace_lock  0   40        0      0 0.00     0        0      0      0
k0x00657e40  procfs_pr_trace_lock  1   40        0      0 0.00     0        0      0      0
:
:
```

2.2.3.15 Displaying lockinfo Structures

The `lockinfo` extension displays static lock class information contained in the `lockinfo` structures. Each lock class is recorded in one `lockinfo` structure, which contains the following information:

- The address of the structure
- The index into the array of `lockinfo` structures
- The class of lock for which information is provided
- The number of instances of the lock
- The lock flag, as defined in the `/sys/include/sys/lock.h` header file

This extension is available only when the `lockmode` system attribute is set to 4.

This extension has the following format:

lockinfo [-class *name*]

The `-class` flag allows you to display the `lockinfo` structure for a particular class of locks. If you omit the flag, `lockinfo` displays the `lockinfo` structures for all classes of locks.

For example:

```
(kdbx) lockinfo
Lockinfo      Index      li_name      li_count  li_flgsp
=====
fffffc0000652030    3      cfg_subsys_lock      21      0xd0
fffffc0000652040    4      subsys_tbl_lock      1      0xc0
fffffc0000652050    5      inode.i_io_lock      4348     0x90
fffffc0000652060    6      nfs_daemon_lock      1      0xc0
fffffc0000652070    7      lk_lmf              1      0xc0
fffffc0000652080    8      procfs_global_lock   1      0xc0
fffffc0000652090    9      procfs.pr_trace_lock 40      0xc0
fffffc00006520a0   10      procnode.prc_ioctl_lock 0      0xc0
fffffc00006520b0   11      semidq_lock          1      0xc0
fffffc00006520c0   12      semid_lock           16      0xc0
fffffc00006520d0   13      undo_lock             1      0xc0
fffffc00006520e0   14      msgidq_lock           1      0xc0
fffffc00006520f0   15      msgid_lock            64      0xc0
fffffc0000652100   16      pgrphash_lock         1      0xc0
fffffc0000652110   17      proc_relation_lock    1      0xc0
fffffc0000652120   18      pgrp.pg_lock          20      0xd0
```

2.2.3.16 Displaying the Mount Table

The `mount` extension displays the mount table, and has the following format:

mount [-s] [*address*]

The `-s` flag displays a short form of the table. If you specify one or more addresses, `kdbx` displays the mount entries named by the addresses.

For example:

```
(kdbx) mount
MOUNT      MAJ      MIN      VNODE      ROOTVP      TYPE      PATH      FLAGS
=====
v0x8196bb30    8          0      NULL      v0x8a75f600  ufs      /
loc
v0x8196a910      v0x8a62de00 v0x8a684e00  nfs      /share/cia/build/alpha.dsk5  ro
v0x8196aae0      v0x8a646800 v0x8a625400  nfs      /share/xor/build/agosminor.dsk1  ro
v0x8196acb0      v0x8a684800 v0x8a649400  nfs      /share/buffer/build/submit.dsk2  ro
v0x8196ae80      v0x8a67ea00 v0x8a774800  nfs      /share/cia/build/goldos.dsk6  ro
v0x8196b050      v0x8a67c400 v0x8a767800  nfs      /usr/staff/alpha1/user
v0x8196b220      v0x8a651800 v0x8a781000  nfs      /usr/sde
ro
v0x8196b3f0    8      2050  v0x8a61ca00 v0x8a77fe00  ufs      /usr3
loc
v0x8196b5c0    8          7  v0x8a61c000 v0x8a79c200  ufs      /usr2
loc
v0x8196b790    8          6  v0x8a5c4800 v0x8a760600  ufs      /usr
loc
v0x8196b960    0          0  v0x8a5c5000 NULL          procfs /proc
```

2.2.3.17 Displaying the Namecache Structures

The namecache extension displays the namecache structures on the system, and has the following format:

namecache

For example:

```
(kdbx) namecache
namecache      nc_vp      nc_vpid  nc_nlen  nc_dvp      nc_name
=====
v0x9047b2c0   v0x9021f4f8      24       4   v0x9021e5b8  sbin
v0x9047b310   v0x9021e988       0      11   v0x9021e7a0  swapdefault
v0x9047b360   v0x9021e5b8       0       2   v0x9021e7a0  ..
v0x9047b3b0   v0x9021e7a0     199       3   v0x9021e5b8  dev
v0x9047b400   v0x9021ed58       0       4   v0x9021eb70  rz1g
v0x9047b4a0   v0x9021f128       0       4   v0x9021e7a0  init
v0x9047b4f0   v0x9021f310       0       7   v0x9021e5b8  upgrade
v0x9047b540   v0x9021fab0      20       3   v0x9021e5b8  etc
v0x9047b590   v0x9021f6e0       0       7   v0x9021f4f8  inittab
v0x9047b5e0   v0x9021eb70      28       3   v0x9021e5b8  var
v0x9047b630   v0x9021f310      34       3   v0x9021e5b8  usr
v0x9047b6d0   v0x9021fc98       0       7   v0x9021eb70  console
v0x9047b720   v0x9021fe80       0       2   v0x9021e7a0  sh
v0x9047b770   v0x90220068       0       3   v0x9021f4f8  nls
v0x9047b810   v0x90220250       0       8   v0x9021e7a0  bcheckrc
v0x9047b8b0   v0x90220438       0       4   v0x9021e7a0  fsck
v0x9047b900   v0x90220620       0       5   v0x9021f4f8  fstab
v0x9047b950   v0x90220808       0       8   v0x9021e7a0  ufs_fsck
v0x9047b9a0   v0x902209f0       0       4   v0x9021eb70  rz1a
v0x9047b9f0   v0x90220bd8       0       5   v0x9021eb70  rrz1a
:
:
```

2.2.3.18 Displaying Processes' Open Files

The ofile extension displays the open files of processes and has the following format.

ofile [-proc *address* | -pid *pid* | -v]

If you omit arguments, ofile displays the files opened by each process. If you specify -proc *address* or -pid *pid* the extension displays the open files owned by the specified process. The -v flag displays more information about the open files.

For example:

```
(kdbx) ofile -pid 1136 -v
Proc=0xffffffff9041e980 pid= 1136
ADDR_FILE  f_cnt  ADDR_VNODE  V_TYPE  V_TAG  USECNT  V_MOUNT  INO#  QSIZE
=====
v0x90408520  27    v0x902c1390  VCHR   VT_UFS   3      v0x863abab8  1103  0
v0x90408520  27    v0x902c1390  VCHR   VT_UFS   3      v0x863abab8  1103  0
v0x90408520  27    v0x902c1390  VCHR   VT_UFS   3      v0x863abab8  1103  0
v0x90408368  1     v0x9026e6b8  VDIR   VT_UFS  18     v0x863ab728  64253  512
```

2.2.3.19 Converting the Contents of Memory to Symbols

The `paddr` extension converts a range of memory to symbolic references and has the following format:

paddr *address number-of-longwords*

The arguments to the `paddr` extension are as follows:

address The starting address.

number-of-longwords The number of longwords to display.

For example:

```
(kdbx) paddr 0xffffffff90be36d8 20
[., 0xffffffff90be36d8]: [h_kmem_free_memory_:824, 0xfffffc000037f47c] 0x0000000000000000
[., 0xffffffff90be36e8]: [., 0xfffffffff8b300d30] [hardclock:394, 0xfffffc00002a7d5c]
[., 0xffffffff90be36f8]: 0x0000000000000000 [., 0xfffffffff863828a0]
[., 0xffffffff90be3708]: [setconf:133, 0xfffffc00004949b0] [., 0xffffffff90be39f4]
[., 0xffffffff90be3718]: 0x00000000000004e0 [thread_wakeup_prim:858, 0xfffffc0000328454]
[., 0xffffffff90be3728]: 0x0000000000000001 0xfffffffff000000c
[., 0xffffffff90be3738]: [., 0xffffffff9024e518] [hardclock:394, 0xfffffc00002a7d5c]
[., 0xffffffff90be3748]: 0x0000000004d5ff8 0xffffffffffffffd4
[., 0xffffffff90be3758]: 0x0000000000bc688 [setconf:133, 0xfffffc00004946f0]
[., 0xffffffff90be3768]: [thread_wakeup_prim:901, 0xfffffc00003284d0]
0x000003ff85ef4ca0
```

2.2.3.20 Displaying the Process Control Block for a Thread

The `pcb` extension displays the process control block for a given thread structure located at `thread_address`. The extension also displays the contents of integer and floating-point registers (if nonzero).

This extension has the following format:

pcb *thread_address*

For example:

```
(kdbx) pcb 0xffffffff863a5bc0
Addr pcb          ksp          usp          pc          ps
v0x90e8c000      v0x90e8fb88      0x0          0xfffffc00002dc110  0x5
sp              ptbr          pcb_physaddr
0xffffffff90e8fb88      0x2ad4      0x55aa000
r9      0xffffffff863a5bc0
r10     0xffffffff863867a0
r11     0xffffffff86386790
r13     0x5
```

2.2.3.21 Formatting Command Arguments

The `printf` extension formats one argument at a time to work around the `dbx` debugger's command length limitation. It also supports the `%s` string

substitution, which the dbx debugger's printf command does not. This extension has the following format:

printf *format-string* [*args*]

The arguments to the printf extension are as follows:

format-string A character string combining literal characters with conversion specifications.

args The arguments for which you want kdbx to display values.

For example:

```
(kdbx) printf "allproc = 0x%x" allproc
allproc = 0xffffffff902356b0
```

2.2.3.22 Displaying the Process Table

The proc extension displays the process table. This extension has the following format:

proc [*address*]

If you specify an address, the proc extension displays only the proc structures at that address; otherwise, the extension displays all proc structures.

For example:

```
(kdbx) proc
:
```

Addr	PID	PPID	PGRP	UID	NICE	SIGCATCH	P_SIG	Event	Flags
v0x8191e210	0	0	0	0	0	00000000	00000000		NULL in sys
v0x8197cd80	1	0	1	0	0	207a7eff	00000000		NULL in pagv exec
v0x8198a210	13	1	13	0	0	00002000	00000000		NULL in pagv
v0x819a8d80	120	1	120	0	0	00086001	00000000		NULL in pagv
v0x819a8210	122	1	122	0	0	00004001	00000000		NULL in pagv
v0x81a14210	5249	1	5267	1138	0	00081000	00000000		NULL in pagv exec
v0x819b6210	131	1	131	0	0	20006003	00000000		NULL in pagv
v0x81a18d80	5266	5267	5267	1138	0	00080000	00000000		NULL in pagv exec
v0x81a2ed80	5267	4938	5267	1138	0	00007efb	00000000		NULL in pagv exec
v0x81a42d80	5268	5266	5267	1138	0	00004007	00000000		NULL in pagv exec
v0x81a18210	5270	5273	5267	1138	0	00000000	00000000		NULL in pagv exec
v0x8198ed80	5273	5266	5267	1138	0	00000000	00000000		NULL in pagv exec
v0x81a0ad80	5276	5279	5276	1138	0	01880003	00000000		NULL
in pagv ctty exec									
v0x81a26d80	5278	5249	5278	1138	0	00080002	00000000		NULL
in pagv ctty exec									
v0x819f2d80	5279	1	5267	1138	0	00081000	00000000		NULL in pagv exec
v0x81a14d80	5281	1	5267	1138	0	00081000	00000000		NULL in pagv exec
v0x81a3cd80	5287	5281	5287	1138	0	01880003	00000000		NULL
in pagv ctty exec									

```

v0x81a28210 5301 5276 5301 1138 0 00080002 00000000 NULL
in pagv ctty exec
v0x819aad80 195 1 195 0 0 00080628 00000000 NULL in pagv
v0x8197c210 6346 1 6346 0 0 00004006 00000000 NULL in pagv exec
v0x819c4210 204 1 0 0 0 00086efe 00000000 NULL in pagv
:

```

2.2.3.23 Converting an Address to a Procedure name

The `procaddr` extension converts the specified address to a procedure name. This extension has the following format:

procaddr [*address*]

For example:

```

(kdbx) procaddr callout.c_func
xpt_pool_free

```

2.2.3.24 Displaying Sockets from the File Table

The `socket` extension displays those files from the file table that are sockets with nonzero reference counts. This extension has the following format:

socket

For example:

```

(kdbx) socket
Fileaddr  Sockaddr  Type      PCB      Qlen Qlim Scc  Rcc
=====  =====  =====  =====  =====  =====  ===  =====
v0x904061b8 v0x863b5c08 DGRAM    v0x8632dc88  0    0    0    0
v0x90406370 v0x863b5a08 DGRAM    v0x8632db08  0    0    0    0
v0x90406478 v0x863b5808 DGRAM    v0x8632da88  0    0    0    0
v0x904064d0 v0x863b5608 DGRAM    v0x8632d688  0    0    0    0
v0x904065d8 v0x863b5408 DGRAM    v0x8632dc08  0    0    0    0
v0x90406630 v0x863b5208 DGRAM    v0x8632d588  0    0    0    0
v0x904067e8 v0x863b4208 DGRAM    v0x8632d608  0    0    0    0
v0x90406840 v0x863b4008 DGRAM    v0x8632d788  0    0    0    0
v0x904069a0 v0x8641f008 STRM     v0x8632c808  0    0    0    0
v0x90406aa8 v0x863b4c08 STRM     v0x8632d508  0    2    0    0
v0x90406bb0 v0x863b4e08 STRM     v0x8632da08  0    0    0    0
:

```

2.2.3.25 Displaying a Summary of the System Information

The `sum` extension displays a summary of system information and has the following format:

sum

For example:

```
(kdbx) sum
Hostname : system.dec.com
cpu: DEC3000 - M500      avail: 1
Boot-time:      Tue Nov  3 15:01:37 1992
Time:   Fri Nov  6 09:59:00 1998
Kernel : OSF1 release 1.2 version 1.2 (alpha)
(kdbx)
```

2.2.3.26 Displaying a Summary of Swap Space

The swap extension displays a summary of swap space and has the following format:

swap

For example:

```
(kdbx) swap
-----
Swap device name      Size      In Use      Free
-----
/dev/rz3b             131072k    32424k      98648k  Dumpdev
                    16384p     4053p       12331p
/dev/rz2b             131072k     8k         131064k
                    16384p     1p          16383p
-----
Total swap partitions:  2      262144k    32432k     229712k
                    32768p     4054p     28714p
(kdbx)
```

2.2.3.27 Displaying the Task Table

The task extension displays the task table. This extension has the following format:

task [*proc_address*]

If you specify addresses, the extension displays the task structures named by the argument addresses; otherwise, the debugger displays all tasks.

For example:

```
(kdbx) task
:
:
Task Addr  Ref  Threads  Map      Swap_state  Utask Addr  Proc Addr  Pid
=====  ==  =====  =====  =====  =====  =====  =====
v0x8191e000 17    15  v0x808f7ef0  INSWAPPED  v0x8191e3b0 v0x8191e210    0
v0x8197cb70  3     1  v0x808f7760  INSWAPPED  v0x8197cf20 v0x8197cd80    1
v0x8198a000  3     1  v0x808f7550  INSWAPPED  v0x8198a3b0 v0x8198a210   13
v0x819a8b70  3     1  v0x808f7340  INSWAPPED  v0x819a8f20 v0x819a8d80   120
v0x819a8000  3     1  v0x808f7290  INSWAPPED  v0x819a83b0 v0x819a8210   122
v0x81a14000  3     1  v0x819f1ad0  INSWAPPED  v0x81a143b0 v0x81a14210  5249
v0x819b6000  3     1  v0x808f6fd0  INSWAPPED  v0x819b63b0 v0x819b6210   131
v0x81a18b70  3     1  v0x819f1a20  INSWAPPED  v0x81a18f20 v0x81a18d80  5266
v0x81a2eb70  3     1  v0x819f1340  INSWAPPED  v0x81a2ef20 v0x81a2ed80  5267
v0x81a42b70  3     1  v0x819f1080  INSWAPPED  v0x81a42f20 v0x81a42d80  5268
v0x81a18000  3     1  v0x819f1970  INSWAPPED  v0x81a183b0 v0x81a18210  5270
v0x8198eb70  3     1  v0x808f74a0  INSWAPPED  v0x8198ef20 v0x8198ed80  5273
```

```

v0x81a0ab70 3 1 v0x819f1ce0 INSWAPPED v0x81a0af20 v0x81a0ad80 5276
v0x81a26b70 3 1 v0x819f1760 INSWAPPED v0x81a26f20 v0x81a26d80 5278
v0x819f2b70 3 1 v0x819f1e40 INSWAPPED v0x819f2f20 v0x819f2d80 5279
v0x81a14b70 3 1 v0x819f1b80 INSWAPPED v0x81a14f20 v0x81a14d80 5281
v0x81a3cb70 3 1 v0x819f11e0 INSWAPPED v0x81a3cf20 v0x81a3cd80 5287
v0x81a28000 3 1 v0x819f1550 INSWAPPED v0x81a283b0 v0x81a28210 5301
v0x819aab70 3 1 v0x808f71e0 INSWAPPED v0x819aaf20 v0x819aad80 195
v0x8197c000 3 1 v0x808f76b0 INSWAPPED v0x8197c3b0 v0x8197c210 6346
v0x819c4000 3 1 v0x808f6e70 INSWAPPED v0x819c43b0 v0x819c4210 204
:

```

2.2.3.28 Displaying Information About Threads

The `thread` extension displays information about threads and has the following format:

thread [*proc_address*]

If you specify addresses, the `thread` extensions displays thread structures named by the addresses; otherwise, information about all threads is displayed.

For example:

```

(kdbx) thread
Thread Addr  Task Addr  Proc Addr  Event          pcb          state
=====
v0x8644d690  v0x8637e440  v0x9041e830  v0x86420668  v0x90f50000  wait
v0x8644d480  v0x8637e1a0  v0x9041eec0  v0x86421068  v0x90f48000  wait
v0x863a17b0  v0x86380ba0  v0x9041db10  v0x8640e468  v0x90f30000  wait
v0x863a19c0  v0x86380e40  v0x9041d9c0  v0x8641f268  v0x90f2c000  wait
v0x8644dccc0  v0x8637ec20  v0x9041e6e0  v0x8641fc00  v0x90f38000  wait
v0x863a0520  v0x8637f400  v0x9041ed70  v0x8640ea00  v0x90f3c000  wait
v0x863a0310  v0x8637f160  v0x9041e980  u0x00000000  v0x90f44000  run
v0x863a2410  v0x863818c0  v0x9041dc60  v0x8640f268  v0x90f18000  wait
v0x863a15a0  v0x86380900  v0x9041d480  v0x8641ec00  v0x90f24000  wait
:
:

```

2.2.3.29 Displaying a Stack Trace of Threads

The `trace` extension displays the stack of one or more threads. This extension has the following format:

trace [*thread_address...* | `-k` | `-u` | `-a`]

If you omit arguments, `trace` displays the stack trace of all threads. If you specify a list of thread addresses, the debugger displays the stack trace of the specified threads. The following table explains the `trace` flags:

- `-a` Displays the stack trace of the active thread on each CPU
- `-k` Displays the stack trace of all kernel threads

-u Displays the stack trace of all user threads

For example:

```
(kdbx) trace
*** stack trace of thread 0xffffffff819af590 pid=0 ***
> 0 thread_run(new_thread = 0xffffffff819af928)
["../../../../src/kernel/kern/sched_prim.c":1637, 0xfffffc00002f9368]
  1 idle_thread() ["../../../../src/kernel/kern/sched_prim.c":2717,
0xfffffc00002fa32c]
*** stack trace of thread 0xffffffff819af1f8 pid=0 ***
> 0 thread_block() ["../../../../src/kernel/kern/sched_prim.c":1455,
0xfffffc00002f9084]
  1 softclock_main() ["../../../../src/kernel/bsd/kern_clock.c":810,
0xfffffc000023a6d4]
:
:

*** stack trace of thread 0xffffffff819fc398 pid=0 ***
> 0 thread_block() ["../../../../src/kernel/kern/sched_prim.c":1471,
0xfffffc00002f9118]
  1 vm_pageout_loop() ["../../../../src/kernel/vm/vm_pagelru.c":375,
0xfffffc0000395664]
  2 vm_pageout() ["../../../../src/kernel/vm/vm_pagelru.c":834,
0xfffffc00003961e0]
:
:

*** stack trace of thread 0xffffffff819fce60 pid=2 ***
> 0 thread_block() ["../../../../src/kernel/kern/sched_prim.c":1471,
0xfffffc00002f9118]
  1 msg_dequeue(message_queue = 0xffffffff819a5970, max_size = 8192,
option = 0, tout = 0, kmsgptr = 0xffffffff916e3980)
["../../../../src/kernel/kern/ipc_basics.c":884, 0xfffffc00002e8b54]
  2 msg_receive_trap(header = 0xfffffc00005bc150, option = 0, size =
8192, name = 0, tout = 0)
["../../../../src/kernel/kern/ipc_basics.c":1245, 0xfffffc00002e92a4]
  3 msg_receive(header = 0xfffffc00005be150, option = 6186352, tout =
0) ["../../../../src/kernel/kern/ipc_basics.c":1107, 0xfffffc00002e904c]
  4 ux_handler() ["../../../../src/kernel/builtin/ux_exception.c":221,
0xfffffc000027269c]
*** stack trace of thread 0xffffffff81a10730 pid=13 ***
> 0 thread_block() ["../../../../src/kernel/kern/sched_prim.c":1471,
0xfffffc00002f9118]
  1 mpsleep(chan = 0xffffffff819f3270 =

"H4\237\201\377\377\377\377^X0\237\201\377\377\377\377^ ^YR", pri =
296, wmesg = 0xfffffc000042f5e0 =

"\200B\260\300B\244KA\340\3038F]\244\377, timo = 0,
lockp = (nil), flags = 0)
["../../../../src/kernel/bsd/kern_synch.c":341, 0xfffffc0000250250]
  2 sigsuspend(p = 0xffffffff81a04278, args = 0xffffffff9170b8a8,
retval = 0xffffffff9170b898)
:
:
```

2.2.3.30 Displaying a u Structure

The `u` extension displays a `u` structure. This extension has the following format:

u [*proc-addr*]

If you omit arguments, the extension displays the `u` structure of the currently running process.

For example:

```
(kdbx) u ffffffff9027ff38
procp 0x9027ff38
ar0 0x90c85ef8
comm cfgmgr
args g B* ü
u_ofile_of: 0x86344e30 u_pofile_of: 0x86345030
0 0xffffffff902322d0
1 0xffffffff90232278
2 0xffffffff90232278
3 0xffffffff90232328
4 0xffffffff90232380 Auto-close
5 0xffffffff902324e0
sizes 29 45 2 (clicks)
u_outime 0
sigs
    40 40 40 40 40 40 40 40
    40 40 40 40 40 40 40 40
    40 40 40 40 40 40 40 40
    40 40 40 40 40 40 40 40
sigmask
    0 fffefeff fffefeff fffefeff 0 0 0 0
    0 0 0 0 0 fffefeff 0 fffefeff
    0 0 0 0 0 0 0 0
    0 0 0 0 0 0 0 0
sigonstack 0
oldmask 2000
sigstack 0 0
cdir rdir 901885b8 0
timers
start 0 723497702
acflag 193248
(kdbx)
```

2.2.3.31 Displaying References to the `ucred` Structure

The `ucred` extension displays all instances of references to `ucred` structures. This extension has the following format:

ucred [-*proc* | -*uthread* | -*file* | -*buf* | -*ref addr* | -*check addr* | *checkall*]

If you omit all flags, `ucred` displays all references to `ucred` structures. The following describes the flags you can specify:

`-proc` Displays all `ucreds` referenced by the `proc` structures

- uthread Displays all ucreds referenced by the uthread structures
- file Displays all ucreds referenced by the file structures
- buf Displays all ucreds referenced by the buf structures
- ref *address* Displays all references to a given ucred
- check *address* Checks the reference count of a particular ucred
- checkall Checks the reference count of all ucreds, with mismatches marked by an asterisk (*)

For example:

```
(kdbx) ucred
  ADDR OF UCRED      ADDR OF Ref      Ref Type cr_ref cr_uid cr_gid cr_ruid
  =====
0xffffffff863d4960  0xffffffff90420f90   proc     3      0      1      0
0xffffffff8651fb80  0xffffffff9041e050   proc    18      0      1      0
0xffffffff86525c20  0xffffffff90420270   proc     2      0      1      0
0xffffffff86457ea0  0xffffffff90421380   proc     4     1139     15     1139
0xffffffff86457ea0  0xffffffff9041f6a0   proc     4     1139     15     1139
0xffffffff8651b5e0  0xffffffff9041f010   proc     2      0      1      0
0xffffffff8651efa0  0xffffffff9041e1a0   proc     2     1138     10     1138
:
:
0xffffffff863d4960  0xffffffff90fb82e0  uthread   3      0      1      0
0xffffffff8651fb80  0xffffffff90fbc2e0  uthread  18      0      1      0
0xffffffff86525c20  0xffffffff90fb02e0  uthread   2      0      1      0
0xffffffff86457ea0  0xffffffff90f882e0  uthread   4     1139     15     1139
0xffffffff86457ea0  0xffffffff90f902e0  uthread   4     1139     15     1139
0xffffffff8651b5e0  0xffffffff90fc02e0  uthread   2      0      1      0
0xffffffff8651efa0  0xffffffff90fac2e0  uthread   2     1138     10     1138
:
:
0xffffffff863d5c20  0xffffffff90406790   file    16      0      0      0
0xffffffff863d5b80  0xffffffff904067e8   file     7      0      0      0
0xffffffff863d5c20  0xffffffff90406840   file    16      0      0      0
0xffffffff863d5b80  0xffffffff90406898   file     7      0      0      0
0xffffffff86456000  0xffffffff904068f0   file    15     1139     15     1139
0xffffffff863d5c20  0xffffffff90406948   file    16      0      0      0
:
:
(kdbx) ucred -ref 0xffffffff863d5a40
  ADDR OF UCRED      ADDR OF Ref      Ref Type cr_ref cr_uid cr_gid cr_ruid
  =====
0xffffffff863d5a40  0xffffffff9041c0d0   proc     4      0      0      0
0xffffffff863d5a40  0xffffffff90ebc2e0  uthread   4      0      0      0
0xffffffff863d5a40  0xffffffff90406f78   file     4      0      0      0
0xffffffff863d5a40  0xffffffff90408730   file     4      0      0      0
(kdbx) ucred -check 0xffffffff863d5a40
  ADDR OF UCRED      cr_ref  Found
  =====
0xffffffff863d5a40  4       0
0xffffffff863d5a40  4       0
0xffffffff863d5a40  4       0
0xffffffff863d5a40  4       0
```

```
=====
0xfffffffff863d5a40      4      4
```

2.2.3.32 Removing Aliases

The `unaliasall` extension removes all aliases, including the predefined aliases. This extension has the following format:

unaliasall

For example:

```
(kdbx) unaliasall
```

2.2.3.33 Displaying the vnode Table

The `vnode` extension displays the `vnode` table and has the following format:

```
vnode [-free | -all | -ufs | -nfs | -cdf | -advfs | -fs address | -u uid | -g gid | -v]
```

If you omit flags, `vnode` displays ACTIVE entries in the `vnode` table. (ACTIVE means that `usecount` is nonzero.) The following describes the flags you can specify:

<code>-free</code>	Displays INACTIVE entries in the <code>vnode</code> table
<code>-all</code>	Prints ALL (both ACTIVE and INACTIVE) entries in the <code>vnode</code> table
<code>-ufs</code>	Displays all UFS entries in the <code>vnode</code> table
<code>-nfs</code>	Displays all NFS entries in the <code>vnode</code> table
<code>-cdf</code>	Displays all CDFS entries in the <code>vnode</code> table
<code>-advfs</code>	Displays all ADVFS entries in the <code>vnode</code> table
<code>-fs <i>address</i></code>	Displays the <code>vnode</code> entries of a mounted file system
<code>-u <i>uid</i></code>	Displays <code>vnode</code> entries of a particular user
<code>-g <i>gid</i></code>	Displays <code>vnode</code> entries of a particular group
<code>-v</code>	Displays related <code>inode</code> , <code>rnode</code> , or <code>cdnode</code> information (used with <code>-ufs</code> , <code>-nfs</code> , or <code>-cdf</code> only)

For example:

```
(kdbx) vnode
ADDR_VNODE  V_TYPE  V_TAG  USECNT  V_MOUNT
=====  =====  =====  =====  =====
v0x9021e000  VBLK  VT_NON    1  k0x00467ee8
v0x9021e1e8  VBLK  VT_NON   83  v0x863abab8
v0x9021e3d0  VBLK  VT_NON    1  k0x00467ee8
v0x9021e5b8  VDIR  VT_UFS   34  v0x863abab8
v0x9021e7a0  VDIR  VT_UFS    1  v0x863abab8
v0x9021ed58  VBLK  VT_UFS    1  v0x863abab8
v0x9021ef40  VBLK  VT_NON    1  k0x00467ee8
v0x9021f128  VREG  VT_UFS    3  v0x863abab8
v0x9021f310  VDIR  VT_UFS    1  v0x863abab8
v0x9021f8c8  VREG  VT_UFS    1  v0x863abab8
v0x9021fe80  VREG  VT_UFS    1  v0x863abab8
v0x902209f0  VDIR  VT_UFS    1  v0x863abab8
v0x90220fa8  VBLK  VT_UFS    9  v0x863abab8
v0x90221190  VBLK  VT_NON    1  k0x00467ee8
v0x90221560  VREG  VT_UFS    1  v0x863abab8
v0x90221748  VBLK  VT_UFS  3153  v0x863abab8
:
:

(kdbx) vnode -nfs -v
ADDR_VNODE  V_TYPE  V_TAG  USECNT  V_MOUNT      FILEID  MODE  UID  GID  QSIZE
=====  =====  =====  =====  =====  =====  =====  =====  =====  =====
v0x90246820  VDIR  VT_NFS    1  v0x863ab560  205732  40751  1138  23  2048
v0x902471a8  VDIR  VT_NFS    1  v0x863ab398  378880  40755  1138  10  5120
v0x90247578  VDIR  VT_NFS    1  v0x863ab1d0    2  40755    0  0  1024
v0x90247948  VDIR  VT_NFS    1  v0x863ab008  116736  40755  1114  0  512
v0x9026d1c0  VDIR  VT_NFS    1  v0x863ab1d0  14347  40755    0  10  512
v0x9026e8a0  VDIR  VT_NFS    1  v0x863aae40    2  40755    0  10  512
v0x9026ea88  VDIR  VT_NFS    1  v0x863ab1d0  36874  40755    0  10  512
v0x90272788  VDIR  VT_NFS    1  v0x863ab1d0  67594  40755    0  10  512
v0x902fd080  VREG  VT_NFS    1  v0x863ab1d0  49368  100755  8887  177  455168
v0x902ff888  VREG  VT_NFS    1  v0x863ab1d0  49289  100755  8887  177  538200
v0x90326410  VREG  VT_NFS    1  v0x863aae40  294959  100755    3  4  196608
:
:

(kdbx) vnode -ufs -v
ADDR_VNODE  V_TYPE  V_TAG  USECNT  V_MOUNT      INODE#  MODE  UID  GID  QSIZE
=====  =====  =====  =====  =====  =====  =====  =====  =====  =====
v0x9021e5b8  VDIR  VT_UFS   34  v0x863abab8    2  40755    0  0  1024
v0x9021e7a0  VDIR  VT_UFS    1  v0x863abab8  1088  40755    0  0  2560
v0x9021ed58  VBLK  VT_UFS    1  v0x863abab8  1175  60600    0  0  0
v0x9021f128  VREG  VT_UFS    3  v0x863abab8  7637  100755    3  4  147456
v0x9021f310  VDIR  VT_UFS    1  v0x863abab8  8704  40755    3  4  512
v0x9021f8c8  VREG  VT_UFS    1  v0x863abab8  7638  100755    3  4  90112
v0x9021fe80  VREG  VT_UFS    1  v0x863abab8  7617  100755    3  4  196608
v0x902209f0  VDIR  VT_UFS    1  v0x863abab8  9792  41777    0  10  512
v0x90220fa8  VBLK  VT_UFS    9  v0x863abab8  1165  60600    0  0  0
v0x90221560  VREG  VT_UFS    1  v0x863abab8  7635  100755    3  4  245760
v0x90221748  VBLK  VT_UFS  3151  v0x863abab8  1184  60600    0  0  0
:
:
```


2.3 The kdebug Debugger

The kdebug debugger allows you to debug running kernel programs. You can start and stop kernel execution, examine variable and register values,

and perform other debugging tasks, just as you would when debugging user space programs.

The ability to debug a running kernel is provided through remote debugging. The kernel code you are debugging runs on a test system. The dbx debugger runs on a remote build system. The debugger communicates with the kernel code you are debugging over a serial communication line or through a gateway system. You use a gateway system when you cannot physically connect the test and build systems. Figure 2–1 shows the connections needed when you use a gateway system.

Figure 2–1: Using a Gateway System During Remote Debugging



ZK-0974U-R

As shown in Figure 2–1, when you use a gateway system, the build system is connected to it using a network line. The gateway system is connected to the test system using a serial communication line.

Prior to running the `kdebug` debugger, the test, build, and gateway systems must meet the following requirements:

- The test system must be running Tru64 UNIX Version 2.0 or higher, must have the Kernel Debugging Tools subset loaded, and must have the Kernel Breakpoint Debugger kernel option configured.
- The build system must be running Tru64 UNIX Version 2.0 or higher and must have the Kernel Debugging Tools subset loaded. Also, this system must contain a copy of the kernel code you are testing and, preferably, the source used to build that kernel code.
- The gateway system must be running Tru64 UNIX Version 2.0 or higher and must have the Kernel Debugging Tools subset loaded.

To use the `kdebug` debugger, you must set up your build, gateway, and test systems as described in Section 2.3.1. Once you complete the setup, you invoke `dbx` as described in Section 2.3.2 and enter commands as you normally would. Refer to Section 2.3.3 if you have problems with the setup of your remote `kdebug` debugging session.

2.3.1 Getting Ready to Use the `kdebug` Debugger

To use the `kdebug` debugger, you must do the following:

1. Attach the test system and the build (or gateway) system.

To attach the serial line between the test and build (or gateway) systems, locate the serial line used for kernel debugging. In general, the correct serial line is either `/dev/tty00` or `/dev/tty01`. For example, if you have a DEC 3000 family workstation, `kdebug` debugger input and output is always to the RS232C port on the back of the system. By default, this port is identified as `/dev/tty00` at installation time.

If your system is an AlphaStation or AlphaServer system with an `ace` console serial interface, the system uses one of two serial ports for `kdebug` input and output. By default, these systems use the COMM1 serial port (identified as `/dev/tty00`) when operating as a build or gateway system. These systems use the COMM2 serial port (identified as `/dev/tty01`) when operating as the test system.

To make it easier to connect the build or gateway system and the test system for kernel debugging, you can modify your system setup. You can change the system setup so that the COMM2 serial port is always used for kernel debugging whether the system is operating as a build system, a gateway system, or a test system.

To make COMM2 the serial port used for kernel debugging on AlphaStations and AlphaServers, modify your `/etc/remote` file. On these systems, the default `kdebug` debugger definition in the `/etc/remote` file appears as follows:

```
kdebug:dv=/dev/tty00:br#9600:pa=none:
```

Modify this definition so that the device is `/dev/tty01` (COMM2), as follows:

```
kdebug:dv=/dev/tty01/br#9600:pa=none:
```

2. On the build system, install the Product Authorization Key (PAK) for the Developer's kit (OSF-DEV), if it is not already installed. For the gateway and tests systems, the OSF-BASE license PAK is all that is needed. For information about installing PAKs, see the *Software License Management* guide.
3. On the build system, modify the setting of the `$kdebug_host`, `$kdebug_line`, or `$kdebug_dbgtty` as needed.

The `$kdebug_host` variable is the name of the gateway system. By default, `$kdebug_host` is set to `localhost`, assuming no gateway system is being used.

The `$kdebug_line` variable selects the serial line definition to use in the `/etc/remote` file of the build system (or the gateway system, if one is being used). By default, `$kdebug_line` is set to `kdebug`.

The `$kdebug_dbgtty` variable sets the terminal on the gateway system to display the communication between the build and test systems, which is useful in debugging your setup. To determine the terminal name to supply to the `$kdebug_dbgtty` variable, enter the `tty` command in the correct window on the gateway system. By default, `$kdebug_dbgtty` is null.

For example, the following `$HOME/.dbxinit` file sets the `$kdebug_host` variable to a system named `gateway`:

```
set $kdebug_host="gateway"
```

4. Recompile kernel files, if necessary.

By default, the kernel is compiled with only partial debugging information. Occasionally, this partial information causes `kdebug` to display erroneous arguments or mismatched source lines. To correct this, recompile selected source files on the test system specifying the `CDEBUGOPTS=-g` argument.

5. Make a backup copy of the kernel running on the test system so that you can restore that kernel after testing:

```
# mv /vmunix /vmunix.save
```

6. Copy the kernel to be tested to `/vmunix` on the test system and reboot the system:

```
# cp vmunix.test /vmunix
# shutdown -r now
```

7. If you are debugging on an SMP system, set the `lockmode` system attribute to 4 on the test system, as follows:

- a. Create a stanza-formatted file named, for example `generic.stanza`, that appears as follows:

```
generic:
    lockmode = 4
```

This file indicates that you are modifying the `lockmode` attribute in the `generic` subsystem.

- b. Use the `sysconfigdb` command to add the contents of the file to the `/etc/sysconfigtab` database:

```
# sysconfigdb -a -f generic.stanza generic
```

- c. Reboot your system.

Setting this system attribute makes debugging on an SMP system easier. For information about the advantages provided see Section 2.1.11.

8. Set the `OPTIONS KDEBUG` configuration file option in your test kernel. To set this option, run the `doconfig` command without flags, as shown:

```
# doconfig
```

Choose `KERNEL BREAKPOINT DEBUGGING` from the kernel options menu when it is displayed by `doconfig`. Once `doconfig` finishes building a new kernel, copy that kernel to the `/vmunix` file and reboot your system. For more information about using the kernel options menu to modify the kernel, see the *System Administration* manual.

2.3.2 Invoking the kdebug Debugger

You invoke the `kdebug` debugger as follows:

1. Invoke the `dbx` debugger on the build system, supplying the pathname of the test kernel. Set a breakpoint and start running `dbx` as follows:

```
# dbx -remote vmunix
dbx version 5.0
Type 'help' for help.
main: 602 p = &proc[0];
(dbx) stop in main
[2] stop in main
(dbx) run
```

Note that you can set a breakpoint anytime after the execution of the `kdebug_bootstrap()` routine. Setting a breakpoint prior to the execution of this routine can result in unpredictable behavior.

You can use all valid `dbx` flags with the `-remote` flag and define entries in your `$HOME/.dbxinit` file as usual. For example, suppose you start the `dbx` session in a directory other than the one that contains the source and object files used to build the `vmunix` kernel you are running on the test system. In this case, use the `-I` command flag or the `use` command in your `$HOME/.dbxinit` file to point `dbx` to the appropriate source and object files. For more information, see `dbx(1)` and the *Programmer's Guide*.

2. Halt the test system and, at the console prompt (three right angle brackets), set the `boot_osflags` console variable to contain the `k` option, and then boot the system. For example:

```
>>> set boot_osflags "k"
>>> boot
```

Once you boot the kernel, it begins executing. The `dbx` debugger will halt execution at the breakpoint you specified, and you can begin issuing

dbx debugging commands. See Section 2.1, the dbx(1) reference page, or the *Programmer's Guide* for information on dbx debugging commands.

If you are unable to bring your test kernel up to a fully operational mode, you can reboot the halted system running the generic kernel, as follows:

```
>>> set boot_osflags "S"
>>> set boot_file "/genvmunix"
>>> boot
```

Once the system is running, you can run the `bcheckrc` script manually to check and mount your local file systems. Then, copy the appropriate kernel to the `root (/)` directory.

When you are ready to resume debugging, copy the test kernel to `/vmunix` and reset the console variables and boot the system, as follows:

```
>>> set boot_osflags "k"
>>> set boot_file "/vmunix"
>>> boot
```

When you have completed your debugging session, reset the console variables on the test system to their normal values, as follows:

```
>>> set boot_osflags "A"
>>> set boot_file "/vmunix"
>>> set auto_action boot
```

You might also need to replace the test kernel with a more reliable kernel. For example, you should have saved a copy of the `vmunix` file that is normally used to run the test system. You can copy that file to `/vmunix` and shut down and reboot the system:

```
# mv /vmunix.save /vmunix
# shutdown -r now
```

2.3.3 Diagnosing kdebug Setup Problems

If you have completed the `kdebug` setup as described in Section 2.3.2 and it fails to work, refer to the following list for help in diagnosing and fixing the setup problem:

- Determine whether the serial line is attached properly and then use the `tip` command to test the connection.

Once you determine that the serial line is attached properly, log on to the build system (or the gateway system if one is being used) as `root` and enter the following command:

```
# tip kdebug
```

If the command does not return the message `connected`, another process, such as a print daemon, might be using the serial line port that you have dedicated to the `kdebug` debugger. To remedy this condition, do the following:

- Check the `/etc/inittab` file to see if any processes are using that line. If so, disable these lines until you finish with the `kdebug` session. See the `inittab(4)` reference page for information on disabling lines.
- Examine your `/etc/remote` file to determine which serial line is associated with the `kdebug` label. Then, use the `ps` command to see if any processes are using the line. For example, if you are using the `/dev/tty00` serial port for your `kdebug` session, check for other processes using the serial line with the following command:

```
# ps agxt00
```

If a process is using `tty00`, either kill that process or modify the `kdebug` label so that a different serial line is used.

If the serial line specified in your `/etc/remote` file is used as the system's serial console, do not kill the process. In this case, use another serial line for the `kdebug` debugger.

- Determine whether any unused `kdebugd` gateway daemons are running with the following command:

```
# ps agx | grep kdebugd
```

After ensuring the daemons are unused, kill the daemon processes.

- If the test system boots to single user or beyond, then `kdebug` has not been configured into the kernel as specified in Section 2.3.1. Ensure that the `boot_osflags` console environment variable specifies the `k` flag and try booting the system again:

```
>>> set boot_osflags k
>>> boot
```

- Be sure you defined the `dbx` variables in your `$HOME/.dbxinit` file correctly.

Determine which terminal line you ran `tip` from by issuing the `/usr/bin/tty` command. For example:

```
# /usr/bin/tty
/dev/ttyp2
```

This example shows that you are using terminal `/dev/ttyp2`. Edit your `$HOME/.dbxinit` file on the build system as follows:

- Set the `$kdebug_dbgtty` variable to `/dev/ttyp2` as follows:

```
set $kdebug_dbgtty="/dev/ttyp2"
```

- Set the `$kdebug_host` variable to the host name of the system from which you entered the `tip` command. For example, if the host name is `MYSYS`, the entry in the `$HOME/.dbxinit` file will be as follows:

```
set $kdebug_host="mysys"
```

- Remove any settings of the `$kdebug_line` variable as follows:

```
set $kdebug_line=
```

- Start `dbx` on the build system. You should see informational messages on the terminal line `/dev/tty2` that `kdebug` is starting.
- If you are using a gateway system, ensure that the `inetd` daemon is running on the gateway system. Also, check the TCP/IP connection between the build and gateway systems using one of the following commands: `rlogin`, `rsh`, or `rcp`.

2.3.4 Notes on Using the `kdebug` Debugger

The following list contains information that can help you use the `kdebug` debugger effectively:

- Breakpoint behavior on SMP systems

If you set breakpoints in code that is executed on an SMP system, the breakpoints are handled serially. When a breakpoint is encountered on a particular CPU, the state of all the other processors in the system is saved and those processors spin. This behavior is similar to how execution stops when a simple lock is obtained on a particular CPU.

Processing resumes on all processors when the breakpoint is dismissed; for example, when you enter a `step` or `cont` command to the debugger.

- Reading instructions from disk

By default, the `dbx` debugger reads instructions from the remote kernel's memory. Reading instructions from memory allows the debugger to help you examine self-modifying code, such as `spl` routines.

You can force the debugger to look at instructions in the on-disk copy of the kernel by adding the following line to your `$HOME/.dbxinit` file:

```
set $readtextfile = 1
```

Setting the `$readtextfile` variable might improve the speed of the debugger while it is reading instructions.

Be aware that the instructions the debugger reads from the on-disk copy of the kernel might be made obsolete by self-modifying code. The on-disk copy of the kernel does not contain any modifications made to the code as it is running. Obsolete instructions that the debugger reads from the on-disk copy can cause the kernel to fail in an unpredictable way.

2.4 The `crashdc` Utility

The `crashdc` utility collects critical data from operating system crash dump files (`vmzcore.n` or `vmcore.n`) or from a running kernel. You can use the data it collects to analyze the cause of a system crash. The `crashdc` utility

uses existing system tools and utilities to extract information from crash dumps. The information garnered from crash dump files or from the running kernel includes the hardware and software configuration, current processes, the panic string (if any), and swap information.

The `crashdc` utility is invoked each time the system is booted. If it finds a current crash dump, `crashdc` creates a data collection file with the same numerical file name extension as the crash dump (see Section 2.1.1 for information about crash dump names).

You can also invoke `crashdc` manually. The syntax for compressed and uncompressed crash dump files, respectively, is as follows:

```
/bin/crashdc vmunix.n vmzcore.n
```

```
/bin/crashdc vmunix.n vmcore.n
```

See Appendix A for an example of the output from the `crashdc` command.

Writing Extensions to the kdbx Debugger

To assist in debugging kernel code, you can write an extension to the kdbx debugger. Extensions interact with kdbx and enable you to examine kernel data relevant to debugging the source program. This chapter provides the following:

- A list of considerations before you begin writing extensions (Section 3.1)
- A description of the kdbx library routines that you can use to write extensions (Section 3.2)
- Examples of kdbx extensions (Section 3.3)
- Instructions for compiling extensions (Section 3.4)
- Information to help you debug your kdbx extensions (Section 3.5)

The Tru64 UNIX Kernel Debugging Tools subset must be installed on your system before you can create custom extensions to the kdbx debugger. This subset contains header files and libraries needed for building kdbx extensions. See Section 3.1 for more information.

3.1 Basic Considerations for Writing Extensions

Before writing an extension, consider the following:

- The information that is needed
Identify the kernel variables and symbols that you need to examine.
- The means for displaying the information
Display the information so that anyone who needs to use it can read and understand it.
- The need to provide useful error checking
As with any good program, it is important to provide informational error messages in the extension.

Before you write an extension, become familiar with the library routines in the `libkdbx.a` library. These library routines provide convenient methods of extracting and displaying kernel data. The routines are declared in the `/usr/include/kdbx.h` header file and described in Section 3.2.

You should also study the extensions that are provided on your system in the `/var/kdbx` directory. These extensions and the example extensions

discussed in Section 3.3 can help you understand what is involved in writing an extension and provide good examples of using the `kdbx` library functions.

3.2 Standard `kdbx` Library Functions

The `kdbx` debugger provides a number of library functions that are used by the resident extensions. You can use these functions, which are declared in the `./usr/include/kdbx.h` header file, to develop customized extensions for your application. To use the functions, you must include the `./usr/include/kdbx.h` header file in your extension.

The sections that follow describe the special data types defined for use in `kdbx` extensions and the library routines you use in extensions. The library routine descriptions show the routine syntax and describe the routine arguments. Examples included in the descriptions show significant lines in boldface type.

3.2.1 Special `kdbx` Extension Data Types

The routines described in this section use the following special data types: `StatusType`, `Status`, `FieldRec`, and `DataStruct`. The uses of these data types are as follows:

- The `StatusType` data type is used to declare the status type and can take on any one of the following values:
 - `OK`, which indicates that no error occurred
 - `Comm`, which indicates a communication error
 - `Local`, which indicates other types of errors

The following is the type definition for the `StatusType` data type:

```
typedef enum { OK, Comm, Local } StatusType;
```

- The `Status` data type is returned by some library routines to inform the caller of the status of the call. Library routines using this data type fill in the `type` field with the call status from `StatusType`. Upon return, callers check the `type` field, and if it is not set to `OK`, they can pass the `Status` structure to the `print_status` routine to generate a detailed error message.

The following is the type definition for the `Status` data type:

```
typedef struct {
    StatusType type;
    union {
        int comm;
        int local;
    } u;
} Status;
```


The values in `comm` and `local` provide the error code interpreted by `print_status`.

- The `FieldRec` data type, which is used to declare a field of interest in a data structure.

The following is the type definition for the `FieldRec` data type:

```
typedef struct {
    char *name;
    int type;
    caddr_t data;
    char *error;
} FieldRec;
```

The `char *name` declaration is the name of the field in question. The `int type` declaration is the type of the field, for example, `NUMBER`, `STRUCTURE`, `POINTER`. The `caddr_t data` and `char *error` declarations are initially set to `NULL`. The `read_field_vals` function fills in these values.

- The `DataStruct`, data type, which is used to declare data structures with opaque data types.

The following is the type definition for the `DataStruct` data type:

```
typedef long DataStruct;
```

3.2.2 Converting an Address to a Procedure Name

The `addr_to_proc` function returns the name of the procedure that begins the address you pass to the function. If the address is not the beginning of a procedure, then a string representation of the address is returned. The return value is dynamically allocated by `malloc` and should be freed by the extension when it is no longer needed.

This function has the following syntax:

```
char * addr_to_proc(
    long addr);
```

Argument	Input/Output	Description
<i>addr</i>	Input	Specifies the address that you want converted to a procedure name

For example:

```
conf1 = addr_to_proc((long) bus_fields[3].data);
conf2 = addr_to_proc((long) bus_fields[4].data);
sprintf(buf, "Config 1 - %sConfig 2 - %s", conf1, conf2);
free(conf1);
free(conf2);
```

3.2.3 Getting a Representation of an Array Element

The `array_element` function returns a representation of one element of an array. The function returns non-NULL if it succeeds or NULL if an error occurs. When the value of `error` is non-NULL, the `error` argument is set to point to the error message. This function has the following syntax:

```
DataStruct array_element(  
    DataStruct sym,  
    int i,  
    char** error);
```

Argument	Input/Output	Description
<i>sym</i>	Input	Names the array
<i>i</i>	Input	Specifies the index of the element
<i>error</i>	Output	Returns a pointer to an error message, if the return value is NULL

You usually use the `array_element` function with the `read_field_vals` function. You use the `array_element` function to get a representation of an array element that is a structure or pointer to a structure. You then pass this representation to the `read_field_vals` function to get the values of fields inside the structure. For an example of how this is done, see Example 3–4 in Section 3.3.

The first argument of the `array_element` function is usually the result returned from the `read_sym` function.

Note

The `read_sym`, `array_element`, and `read_field_vals` functions are often used together to retrieve the values of an array of structures pointed to by a global pointer. (For more information about using these functions, see the description of the `read_sym` function in Section 3.2.27.)

For example:

```
if((ele = array_element(sz_softc, cntrl, &error)) == NULL){  
    fprintf(stderr, "Couldn't get %d'th element of sz_softc:\n, cntrl");  
    fprintf(stderr, "%s\n", error);  
}
```

3.2.4 Retrieving an Array Element Value

The `array_element_val` function returns the value of an array element. It returns the integer value if the data type of the array element is an integer

data type. It returns the pointer value if the data type of the array element is a pointer data type.

This function returns TRUE if it is successful, FALSE otherwise. When the return value is FALSE, an error message is returned in an argument to the function.

This function has the following syntax:

```
Boolean array_element_val(  
    DataStruct sym,  
    int i,  
    long* ele_ret,  
    char ** error);
```

Argument	Input/Output	Description
<i>sym</i>	Input	Names the array
<i>i</i>	Input	Specifies the index of the element
<i>ele_ret</i>	Output	Returns the value of the pointer
<i>error</i>	Output	Returns a pointer to an error message if the return value is FALSE

You use the `array_element_val` function when the array element is of a basic C type. You also use this function if the array element is of a pointer type and the pointer value is what you actually want. This function returns a printable value. The first argument of the `array_element_val` function usually comes from the returned result of the `read_sym` function.

For example:

```
static char get_ele(array, i)  
DataStruct array;  
int i;  
{  
    char *error, ret;  
    long val;  
  
    if(!array_element_val(array, i, &val, &error)){  
        fprintf(stderr, "Couldn't read array element:\n");  
        fprintf(stderr, "%s\n", error);  
        quit(1);  
    }  
    ret = val;  
    return(ret);  
}
```

3.2.5 Returning the Size of an Array

The `array_size` function returns the size of the specified array. This function has the following syntax:

```
unsigned int array_size(
    DataStruct sym,
    char** error);
```

Argument	Input/Output	Description
<i>sym</i>	Input	Names the array
<i>error</i>	Output	Returns a pointer to an error message if the return value is non-NULL

For example:

```
busses = read_sym("bus_list");

if((n = array_size(busses, &error)) == -1){
    fprintf(stderr, "Couldn't call array_size:\n");
    fprintf(stderr, "%s\n", error);
    quit(1);
}
```

3.2.6 Casting a Pointer to a Data Structure

The `cast` function casts the pointer to a structure as a structure data type and returns the structure. This function has the following syntax:

```
Boolean cast(
    long addr,
    char* type,
    DataStruct* ret_type,
    char** error);
```

Argument	Input/Output	Description
<i>addr</i>	Input	Specifies the address of the data structure you want returned
<i>type</i>	Input	Specifies the datatype of the data structure
<i>ret_type</i>	Output	Returns the name of the data structure
<i>error</i>	Output	Returns a pointer to an error message if the return value is FALSE

You usually use the `cast` function with the `read_field_vals` function. Given the address of a structure, you call the `cast` function to convert the pointer from the type `long` to the type `DataStruct`. Then, you pass the result to the `read_field_vals` function, as its first argument, to retrieve the values of data fields in the structure.

For example:

```
if(!cast(addr, "struct file", &fil, &error)){
    fprintf(stderr, "Couldn't cast address to a file:\n");
    fprintf(stderr, "%s\n", error);
    quit(1);
}
```

3.2.7 Checking Arguments Passed to an Extension

The `check_args` function checks the arguments passed to an extension or displays a help message. The function displays a help message when the user specifies the `-help` flag on the command line.

This function has the following syntax:

```
void check_args(
    int argc,
    char** argv,
    char* help_string);
```

Argument	Input/Output	Description
<i>argc</i>	Input	Passes in the first argument to the command
<i>argv</i>	Input	Passes in the second argument to the command
<i>help_string</i>	Input	Specifies the help message to be displayed to the user

You should include the `check_args` function early in your extension to be sure that arguments are correct.

For example:

```
check_args(argc, argv, help_string);
if(!check_fields("struct sz_softc", fields, NUM_FIELDS, NULL)){
    field_errors(fields, NUM_FIELDS);
    quit(1);
}
```

3.2.8 Checking the Fields in a Structure

The `check_fields` function verifies that the specified function consists of the expected number of fields and that those fields have the correct data type. If the function is successful, `TRUE` is returned; otherwise, the error parts of the affected fields are filled in with errors, and `FALSE` is returned.

This function has the following syntax:

```
Boolean check_fields(
    char* symbol,
    FieldRec* fields,
    int nfields,
```

```
char** hints);
```

Argument	Input/Output	Description
<i>symbol</i>	Input	Names the structure to be checked
<i>fields</i>	Input	Describes the fields to be checked
<i>nfields</i>	Input	Specifies the size of the <i>fields</i> argument
<i>hints</i>	Input	Unused and should always be set to NULL

You should check the structure type using the `check_fields` function before using the `read_field_vals` function to read field values.

For example:

```
FieldRec fields[] = {
    { ".sc_sysid", NUMBER, NULL, NULL },
    { ".sc_aipfts", NUMBER, NULL, NULL },
    { ".sc_lostarb", NUMBER, NULL, NULL },
    { ".sc_lastid", NUMBER, NULL, NULL },
    { ".sc_active", NUMBER, NULL, NULL }
};

check_args(argc, argv, help_string);
if(!check_fields("struct sz_softc", fields, NUM_FIELDS, NULL)){
    field_errors(fields, NUM_FIELDS);
    quit(1);
}
```

3.2.9 Setting the kdbx Context

The `context` function sets the context to user context or proc context. If the context is set to the user context, aliases defined in the extension affect user aliases.

This function has the following syntax:

```
void context(
    Boolean user);
```

Argument	Input/Output	Description
<i>user</i>	Input	Sets the context to user if TRUE or proc if FALSE

For example:

```
if(head) print(head);
context(True);
for(i=0;i<len;i++){
    :
}
```

3.2.10 Passing Commands to the dbx Debugger

The `dbx` function passes a command to the `dbx` debugger. The function has an argument, `expect_output`, that controls when it returns. If you set the `expect_output` argument to `TRUE`, the function returns after the command is sent, and expects the extension to read the output from `dbx`. If you set the `expect_output` argument to `FALSE`, the function waits for the command to complete execution, reads the acknowledgement from `kdbx`, and then returns.

```
void dbx(  
    char* command,  
    Boolean expect_output);
```

Argument	Input/Output	Description
<code>command</code>	Input	Specifies the command to be passed to <code>dbx</code>
<code>expect_output</code>	Input	Indicates whether the extension expects output and determines when the function returns

For example:

```
dbx(out, True);  
if((buf = read_response(&status)) == NULL){  
    print_status("main", &status);  
    quit(1);  
}  
else {  
    process_buf(buf);  
    quit(0);  
}
```

3.2.11 Dereferencing a Pointer

The `deref_pointer` function returns a representation of the object pointed to by a pointer. The function displays an error message if the `data` argument passed is not a valid address.

This function has the following syntax:

```
DataStruct deref_pointer(  
    DataStruct data);
```

Argument	Input/Output	Description
<code>data</code>	Input	Names the data structure that is being dereferenced

For example:

```
structure = deref_pointer(struct_pointer);
```

3.2.12 Displaying the Error Messages Stored in Fields

The `field_errors` function displays the error messages stored in fields by the `check_fields` function. This function has the following syntax:

```
void field_errors(  
    FieldRec* fields,  
    int nfields);
```

Argument	Input/Output	Description
<i>fields</i>	Input	Names the fields that contain the error messages
<i>nfields</i>	Input	Specifies the size of the <i>fields</i> argument

For example:

```
if(!read_field_vals(proc, fields, NUM_FIELDS)){  
    field_errors(fields, NUM_FIELDS);  
    return(False);  
}
```

3.2.13 Converting a Long Address to a String Address

The `format_addr` function converts a 64-bit address of type `long` into a 32-bit address of type `char`. This function has the following syntax:

```
extern char* format_addr(  
    long addr,  
    char* buffer);
```

Argument	Input/Output	Description
<i>addr</i>	Input	Specifies the address to be converted
<i>buffer</i>	Output	Returns the converted address and must be at least 12 characters long

Use this function to save space on the output line. For example, the 64-bit address `0xffffffff12345678` is converted into `v0x12345678`.

For example:

```
static Boolean prfile(DataStruct ele, long vn_addr, long socket_addr)  
{  
    char *error, op_buf[12], *ops, buf[256], address[12], cred[12], data[12];  
    if(!read_field_vals(ele, fields, NUM_FIELDS)){  
        field_errors(fields, NUM_FIELDS);  
        return(False);  
    }  
    if((long) fields[1].data == 0) return(True);  
    if((long) (fields[5].data) == 0) ops = " *Null* ";  
    else if((long) (fields[5].data) == vn_addr) ops = " vnops ";  
    else if((long) (fields[5].data) == socket_addr) ops = " socketops ";  
    else format_addr((long) fields[5].data, op_buf);
```



```

format_addr((long) struct_addr(ele), address);
format_addr((long) fields[2].data, cred);
format_addr((long) fields[3].data, data);
sprintf(buf, "%s %s %4d %4d %s %s %s %6d %s%s%s%s%s%s%s%s",
        address, get_type((int) fields[0].data), fields[1].data,
        fields[2].data, ops, cred, data, fields[6].data,
        ((long) fields[7].data) & FREAD ? " read" : ,
        ((long) fields[7].data) & FWRITE ? " write" : ,
        ((long) fields[7].data) & FAPPEND ? " append" : ,
        ((long) fields[7].data) & FNDELAY ? " ndelay" : ,
        ((long) fields[7].data) & FMARK ? " mark" : ,
        ((long) fields[7].data) & FDEFER ? " defer" : ,
        ((long) fields[7].data) & FASYNC ? " async" : ,
        ((long) fields[7].data) & FSHLOCK ? " shlck" : ,
        ((long) fields[7].data) & FEXLOCK ? " exlck" : );
print(buf);
return(True);
}

```

3.2.14 Freeing Memory

The `free_sym` function releases the memory held by a specified symbol. This function has the following syntax:

```
void free_sym(
    DataStruct sym);
```

Argument	Input/Output	Description
<i>sym</i>	Input	Names the symbol that is using memory that can be freed

For example:

```
free_sym(rec->data);
```

3.2.15 Passing Commands to the kdbx Debugger

The `krash` function passes a command to `kdbx` for execution. You specify the command you want passed to `kdbx` as the first argument to the `krash` function. The second argument allows you to pass quotation marks ("), apostrophes ('), and backslash characters (\) to `kdbx`. The function has an argument, `expect_output`, which controls when it returns. If you set the `expect_output` argument to `TRUE`, the function returns after the command is sent, and expects the extension to read the output from `dbx`. If you set the `expect_output` argument to `FALSE`, the function waits for the command to complete execution, reads the acknowledgement from `kdbx`, and then returns.

This function has the following syntax:

```
void krash(
    char* command,
    Boolean quote,
    Boolean expect_output);
```

Argument	Input/Output	Description
<i>command</i>	Input	Names the command to be executed
<i>quote</i>	Input	If set to TRUE causes the quote character, apostrophe, and backslash to be appropriately quoted so that they are treated normally, instead of as special characters
<i>expect_output</i>	Input	Indicates whether the extension expects output and determines when the function returns

For example:

```
do {
    :
    if(doit){
        format(command, buf, type, addr, last, i, next);
        context(True);

        krash(buf, False, True);
        while((line = read_line(&status)) != NULL){
            print(line);
            free(line);
        }
    }
    :
    addr = next;
    i++;
}
```

Suppose the preceding example is used to list the addresses of each node in the system mount table, which is a linked list. The following list describes the arguments to the `format` function in this case:

- The `command` argument contains the dbx command to be executed, such as `p` for `print`.
- The `buf` argument contains the full dbx command line; for example, `buf` might contain:

```
p ((struct mount *) 0xffffffff8196db30).m_next
```
- The `type` argument contains the data type of each node in the list, as in `struct mount *`.
- The `addr` argument contains the address of the current node in the list; for example, the current node might be at address `0xffffffff8196db30`.
- The `last` argument contains the address of the previous node in the list. In this case, `last` contains zero (0).
- The `i` argument is the current node's index. In this case, `i` contains 1.

- The `next` argument is the address of the next node in the list; for example, the next node might be at address `0xffffffff8196d050`.

3.2.16 Getting the Address of an Item in a Linked List

The `list_nth_cell` function returns the address of one of the items in a linked list. This function has the following format:

```
Boolean list_nth_cell(
    long addr,
    char* type,
    int n,
    char* next_field,
    Boolean do_check,
    long* val_ret,
    char** error);
```

Argument	Input/Output	Description
<code>addr</code>	Input	Specifies the starting address of the linked list
<code>type</code>	Input	Specifies the data type of the item for which you are requesting an address
<code>n</code>	Input	Supplies a number indicating which list item's address is being requested
<code>next_field</code>	Input	Gives the name of the field that points to the next item in the linked list
<code>do_check</code>	Input	Determines whether <code>kdbx</code> checks the arguments to ensure that correct information is being sent (TRUE setting)
<code>val_ret</code>	Output	Returns the address of the requested list item
<code>error</code>	Output	Returns a pointer to an error message if the return value is FALSE

For example:

```
long root_addr, addr;
if (!read_sym_val("rootfs", NUMBER, &root_addr, &error)){
    :
}

if(!list_nth_cell(root_addr, "struct mount", i, "m_next", True, &addr,
    &error)){
    fprintf(stderr, "Couldn't get %d'th element of mount table\n", i);
    fprintf(stderr, "%s\n", error);
    quit(1);
}
```

3.2.17 Passing an Extension to kdbx

The `new_proc` function directs `kdbx` to execute a `proc` command with arguments specified in `args`. The `args` arguments can name an extension that is included with the operating system or an extension that you create.

This function has the following syntax:

```
void new_proc(  
    char* args,  
    char** output_ret);
```

Argument	Input/Output	Description
<code>args</code>	Input	Names the extensions to be passed to <code>kdbx</code>
<code>output_ret</code>	Output	Returns the output from the extension, if it is non-NULL

For example:

```
static void pmap(long addr)  
{  
    char cast_addr[36], buf[256], *resp;  
  
    sprintf(cast_addr, "((struct\ vm_map_t\ *)\ 0x%p)", addr);  
    sprintf(buf, "printf  
        cast_addr);  
    new_proc(buf, &resp);  
    print(resp);  
    free(resp);  
}
```

3.2.18 Getting the Next Token as an Integer

The `next_number` function converts the next token in a buffer to an integer. The function returns `TRUE` if successful, or `FALSE` if there was an error.

This function has the following syntax:

```
Boolean next_number(  
    char* buf,  
    char** next,  
    long* ret);
```

Argument	Input/Output	Description
<code>buf</code>	Input	Names the buffer containing the value to be converted
<code>next</code>	Output	Returns a pointer to the next value in the buffer, if that value is non-NULL
<code>ret</code>	Output	Returns the integer value

For example:

```
resp = read_response_status();
next_number(resp, NULL, &size);
ret->size = size;
```

3.2.19 Getting the Next Token as a String

The `next_token` function returns a pointer to the next token in the specified pointer to a string. A token is a sequence of nonspace characters. This function has the following syntax:

```
char* next_token(
    char* ptr,
    int* len_ret,
    char** next_ret);
```

Argument	Input/Output	Description
<i>ptr</i>	Input	Specifies the name of the pointer
<i>len_ret</i>	Output	Returns the length of the next token, if non-NULL
<i>next_ret</i>	Output	Returns a pointer to the first character after, but not included in the current token, if non-NULL

You use this function to extract words or other tokens from a character string. A common use, as shown in the example that follows, is to extract tokens from a string of numbers. You can then cast the tokens to a numerical data type, such as the `long` data type, and use them as numbers.

For example:

```
static long *parse_memory(char *buf, int offset, int size)
{
    long *buffer, *ret;
    int index, len;
    char *ptr, *token, *next;
    NEW_TYPE(buffer, offset + size, long, long *, "parse_memory");
    ret = buffer;
    index = offset;
    ptr = buf;
    while(index < offset + size){
        if((token = next_token(ptr, &len, &next)) == NULL){
            ret = NULL;
            break;
        }
        ptr = next;
        if(token[len - 1] == ':') continue;
        buffer[index] = strtoul(token, &ptr, 16);
        if(ptr != &token[len]){
            ret = NULL;
            break;
        }
        index++;
    }
    if(ret == NULL) free(buffer);
    return(ret);
}
```

```
}
```

3.2.20 Displaying a Message

The `print` function displays a message on the terminal screen. Because of the input and output redirection done by `kdbx`, all output to `stdout` from a `kdbx` extension goes to `dbx`. As a result, a `kdbx` extension cannot use normal C output functions such as `printf` and `fprintf(stdout, ...)` to display information on the screen. Although the `fprintf(stderr, ...)` function is still available, the recommended method is to first use the `sprintf` function to print the output into a character buffer and then use the `kdbx` library function `print` to display the contents of the buffer to the screen.

The `print` function automatically displays a newline character at the end of the output, it fails if it detects a newline character at the end of the buffer.

This function has the following format:

```
void print(  
    char* message);
```

Argument	Input/Output	Description
<i>message</i>	Input	The message to be displayed

For example:

```
if(do_short){  
    if(!check_fields("struct mount", short_mount_fields,  
                    NUM_SHORT_MOUNT_FIELDS, NULL)){  
        field_errors(short_mount_fields, NUM_SHORT_MOUNT_FIELDS);  
        quit(1);  
    }  
    print("SLOT MAJ MIN TYPE                DEVICE MOUNT POINT");  
}
```

3.2.21 Displaying Status Messages

The `print_status` function displays a status message that you supply and a status message supplied by the system. This function has the following format:

```
void print_status(  
    char* message,  
    Status* status);
```

Argument	Input/Output	Description
<i>message</i>	Input	Specifies the extension-defined status message
<i>status</i>	Input	Specifies the status returned from another library routine

For example:

```

if(status.type != OK){
    print_status("read_line failed", &status);
    quit(1);
}

```

3.2.22 Exiting from an Extension

The `quit` function sends a quit command to `kdbx`. This function has the following format:

```

void quit(
    int i);

```

Argument	Input/Output	Description
<i>i</i>	Input	The status at the time of the exit from the extension

For example:

```

if (!read_sym_val("vm_swap_head", NUMBER, &end, &error)) {
    fprintf(stderr, "Couldn't read vm_swap_head:\n");
    fprintf(stderr, "%s\n", error);
    quit(1);
}

```

3.2.23 Reading the Values in Structure Fields

The `read_field_vals` function reads the value of fields in the specified structure. If this function is successful, then the data parts of the fields are filled in and `TRUE` is returned; otherwise, the error parts of the affected fields are filled in with errors and `FALSE` is returned.

This function has the following format:

```

Boolean read_field_vals(
    DataStruct data,
    FieldRec* fields,
    int nfields);

```

Argument	Input/Output	Description
<i>data</i>	Input	Names the structure that contains the field to be read
<i>fields</i>	Input	Describes the fields to be read
<i>nfields</i>	Input	Contains the size of the field array

For example:

```

if(!read_field_vals(pager, fields, nfields)){
    field_errors(fields, nfields);
}

```

```

    return(False);
}

```

3.2.24 Returning a Line of kdbx Output

The `read_line` function returns the next line of the output from the last `kdbx` command executed. If the end of the output is reached, this function returns `NULL` and a status of `OK`. If the status is something other than `OK` when the function returns `NULL`, an error occurred.

This function has the following format:

```

char* read_line(
    Status* status);

```

Argument	Input/Output	Description
<i>status</i>	Output	Contains the status of the request, which is OK for successful requests

For example:

```

while((line = read_line(&status)) != NULL){
    print(line);
    free(line);
}

```

3.2.25 Reading an Area of Memory

The `read_memory` function reads an area of memory starting at the address you specify and running for the number of bytes you specify. The `read_memory` function returns `TRUE` if successful and `FALSE` if there was an error.

This function has the following format:

```

Boolean read_memory(
    long start_addr,
    int n,
    char* buf,
    char** error);

```

Argument	Input/Output	Description
<i>start_addr</i>	Input	Specifies the starting address for the read
<i>n</i>	Input	Specifies the number of bytes to read
<i>buf</i>	Output	Returns the memory contents
<i>error</i>	Output	Returns a pointer to an error message if the return value is <code>FALSE</code>

You can use this function to look up any type of value; however it is most useful for retrieving the value of pointers that point to other pointers.

For example:

```
start_addr = (long) ((long *)utask_fields[7].data + i-NOFILE_IN_U);
if(!read_memory(start_addr , sizeof(long *), (char *)&val1, &error) ||
!read_memory((long)utask_fields[8].data , sizeof(long *), (char *)&val2,
&error)){
    fprintf(stderr, "Couldn't read_memory\n");
    fprintf(stderr, "%s\n", error);
    quit(1);
}
```

3.2.26 Reading the Response to a kdbx Command

The `read_response` function reads the response to the last kdbx command entered. If any errors occurred, `NULL` is returned and the status argument is filled in.

This function has the following syntax:

```
char* read_response(
    Status* status);
```

Argument	Input/Output	Description
<i>status</i>	Output	Contains the status of the last kdbx command

For example:

```
if(!*argv) Usage();
command = argv;
if(size == 0){
    sprintf(buf, "print sizeof*((%s) 0)", type);
    dbx(buf, True);

    if((resp = read_response(&status)) == NULL){
        print_status("Couldn't read sizeof", &status);
        quit(1);
    }
    size = strtoul(resp, &ptr, 0);
    if(ptr == resp){

        fprintf(stderr, "Couldn't parse sizeof(%s):\n", type);
        quit(1);
    }
    free(resp);
}
```

3.2.27 Reading Symbol Representations

The `read_sym` function returns a representation of the named symbol. This function has the following format:

```
DataStruct read_sym(  
    char* name);
```

Argument	Input/Output	Description
<i>name</i>	Input	Names the symbol, which is normally a pointer to a structure or an array of structures inside the kernel

Often you use the result returned by the `read_sym` function as the input argument of the `array_element`, `array_element_val`, or `read_field_vals` function.

For example:

```
busses = read_sym("bus_list");
```

3.2.28 Reading a Symbol's Address

The `read_sym_addr` function reads the address of the specified symbol. This function has the following format:

```
Boolean read_sym_addr(  
    char* name,  
    long* ret_val,  
    char** error);
```

Argument	Input/Output	Description
<i>name</i>	Input	Names the symbol for which an address is required
<i>ret_val</i>	Output	Returns the address of the symbol
<i>error</i>	Output	Returns a pointer to an error message when the return status is FALSE

For example:

```
if(argc == 0) fil = read_sym("file");  
if(!read_sym_val("nfile", NUMBER, &nfile, &error) ||  
    !read_sym_addr("vnops", &vn_addr, &error) ||  
    !read_sym_addr("socketops", &socket_addr, &error)){  
    fprintf(stderr, "Couldn't read nfile:\n");  
    fprintf(stderr, "%s\n", error);  
    quit(1);  
}
```

3.2.29 Reading the Value of a Symbol

The `read_sym_val` function returns the value of the specified symbol. This function has the following format:

```
Boolean read_sym_val(  
    char* name,  
    int type,  
    long* ret_val,  
    char** error);
```

Argument	Input/Output	Description
<i>name</i>	Input	Names the symbol for which a value is needed
<i>type</i>	Input	Specifies the data type of the symbol
<i>ret_val</i>	Output	Returns the value of the symbol
<i>error</i>	Output	Returns a pointer to an error message when the status is FALSE

You use the `read_sym_val` function to retrieve the value of a global variable. The value returned by the `read_sym_val` function has the type `long`, unlike the value returned by the `read_sym` function which has the type `DataStruct`.

For example:

```
if(argc == 0) fil = read_sym("file");  
if(!read_sym_val("nfile", NUMBER, &file, &error) ||  
    !read_sym_addr("vnops", &vn_addr, &error) ||  
    !read_sym_addr("socketops", &socket_addr, &error)){  
    fprintf(stderr, "Couldn't read nfile:\n");  
    fprintf(stderr, "%s\n", error);  
    quit(1);  
}
```

3.2.30 Getting the Address of a Data Representation

The `struct_addr` function returns the address of a data representation. This function has the following format:

```
char* struct_addr(  
    DataStruct data);
```

Argument	Input/Output	Description
<i>data</i>	Input	Specifies the structure for which an address is needed

For example:

```
if(bus_fields[1].data != 0){  
    sprintf(buf, "Bus #d (0x%p): Name - \"%s\"\\tConnected to - \"%s\",  
            i, struct_addr(bus), bus_fields[1].data, bus_fields[2].data);
```

```

print(buf);
sprintf(buf, "\tConfig 1 - %s\tConfig 2 - %s",
        addr_to_proc((long) bus_fields[3].data),
        addr_to_proc((long) bus_fields[4].data));
print(buf);
if(!prctlr((long) bus_fields[0].data)) quit(1);
print();
}

```

3.2.31 Converting a String to a Number

The `to_number` function converts a string to a number. The function returns `TRUE` if successful, or `FALSE` if conversion was not possible.

This function has the following format:

```

Boolean to_number(
    char* str,
    long* val);

```

Argument	Input/Output	Description
<i>str</i>	Input	Contains the string to be converted
<i>val</i>	Output	Contains the numerical equivalent of the string

This function returns `TRUE` if successful, `FALSE` if conversion was not possible.

For example:

```

check_args(argc, argv, help_string);
if(argc < 5) Usage();
size = 0;
type = argv[1];
if(!to_number(argv[2], &len)) Usage();
addr = strtoul(argv[3], &ptr, 16);
if(*ptr != '\0'){
    if(!read_sym_val(argv[3], NUMBER, &addr, &error)){
        fprintf(stderr, "Couldn't read %s:\n", argv[3]);
        fprintf(stderr, "%s\n", error);
        Usage();
    }
}
}

```

3.3 Examples of kdbx Extensions

This section contains examples of the three types of extensions provided by the `kdbx` debugger:

- Extensions that use lists. Example 3–1 provides a C language template and Example 3–2 is the source code for the `/var/kdbx/callout` extension, which shows how to use linked lists in developing an extension.

- Extensions that use arrays. Example 3–3 provides a C language template and Example 3–4 is the source code for the `/var/kdbx/file` extension, which shows how to develop an extension using arrays.
- Extensions that use global symbols. Example 3–5 is the source code for the `/var/kdbx/sum` extensions, which shows how to pull global symbols from the kernel. A template is not provided because the means for pulling global symbols from a kernel can vary greatly, depending upon the desired output.

Example 3–1: Template Extension Using Lists

```

#include <stdio.h>
#include <kdbx.h>
static char *help_string =

"<Usage info goes here>                \\n\ [1]
";
FieldRec fields[] = {
    { "<name of next field>", NUMBER, NULL, NULL }, [2]
    <data fields>
};

#define NUM_FIELDS (sizeof(fields)/sizeof(fields[0]))

main(argc, argv)
int argc;
char **argv;
{
    DataStruct head;
    unsigned int next;
    char buf[256], *func, *error;

    check_args(argc, argv, help_string);
    if(!check_fields("<name of list structure>", fields, NUM_FIELDS, NULL)){ [3]
        field_errors(fields, NUM_FIELDS);
        quit(1);
    }

    if(!read_sym_val("<name of list head>", NUMBER, (caddr_t *) &next, &error)){ [4]
        fprintf(stderr, "%s\n", error);
        quit(1);
    }
    sprintf(buf, "<table header>"); [5]
    print(buf);
    do {
        if(!cast(next, "<name of list structure>", &head, &error)){ [6]
            fprintf(stderr, "Couldn't cast to a <struct>:\n"); [7]
            fprintf(stderr, "%s:\n", error);
        }
        if(!read_field_vals(head, fields, NUM_FIELDS)){
            field_errors(fields, NUM_FIELDS);
            break;
        }
        <print data in this list cell> [8]
        next = (int) fields[0].data;
    } while(next != 0);
}

```

Example 3–1: Template Extension Using Lists (cont.)

```
    quit(0);  
}
```

- ❶ The help string is output by the `check_args` function if the user enters the `help extension_name` command at the `kdbx` prompt. The first line of the help string should be a one-line description of the extension. The rest should be a complete description of the arguments. Also, each line should end with the string `\\\n\`.
- ❷ Every structure field to be extracted needs an entry. The first field is the name of the next extracted field; the second field is the type. The last two fields are for output and initialize to `NULL`.
- ❸ Specifies the type of the list that is being traversed.
- ❹ Specifies the variable that holds the head of the list.
- ❺ Specifies the table header string.
- ❻ Specifies the type of the list that is being traversed.
- ❼ Specifies the structure type.
- ❽ Extracts, formats, and prints the field information.

Example 3–2: Extension That Uses Linked Lists: `callout.c`

```
#include <stdio.h>  
#include <errno.h>  
#include <kdbx.h>  
  
#define KERNEL  
#include <sys/callout.h>  
  
static char *help_string =  
"callout - print the callout table          \\\n\  
  Usage : callout [cpu]                    \\\n\  
";  
  
FieldRec processor_fields[] = {  
  { ".calltodo.c_u.c_ticks", NUMBER, NULL, NULL },  
  { ".calltodo.c_arg",      NUMBER, NULL, NULL },  
  { ".calltodo.c_func",     NUMBER, NULL, NULL },  
  { ".calltodo.c_next",     NUMBER, NULL, NULL },  
  { ".lbolt",               NUMBER,  NULL, NULL },  
  { ".state",               NUMBER,  NULL, NULL },  
};  
  
FieldRec callout_fields[] = {  
  { ".c_u.c_ticks", NUMBER, NULL, NULL },  
  { ".c_arg",      NUMBER, NULL, NULL },  
  { ".c_func",     NUMBER, NULL, NULL },  
  { ".c_next",    NUMBER, NULL, NULL },  
};
```

Example 3–2: Extension That Uses Linked Lists: callout.c (cont.)

```
#define NUM_PROCESSOR_FIELDS
(sizeof(processor_fields)/sizeof(processor_fields[0]))
#define NUM_CALLOUT_FIELDS (sizeof(callout_fields)/sizeof(callout_fields[0]))

main(int argc, char **argv)
{
    DataStruct processor_ptr, processor, callout;
    long next, ncpus, ptr_val, i;
    char buf[256], *func, *error, arg[13];
    int cpuflag = 0, cpuarg = 0;

    long headptr;
    Status status;
    char *resp;

    if ( !(argc == 1 || argc == 2) ) {
        fprintf(stderr, "Usage: callout [cpu]\n");
        quit(1);
    }

    check_args(argc, argv, help_string);

    if (argc == 2) {
        cpuflag = 1;
        errno = 0;
        cpuarg = atoi(argv[1]);
        if (errno != 0)
            fprintf(stderr, "Invalid argument value for the cpu number.\n");
    }

    if(!check_fields("struct processor", processor_fields, NUM_PROCESSOR_FIELDS,
NULL)){
        field_errors(processor_fields, NUM_PROCESSOR_FIELDS);
        quit(1);
    }

    if(!check_fields("struct callout", callout_fields, NUM_CALLOUT_FIELDS, NULL)){
        field_errors(callout_fields, NUM_CALLOUT_FIELDS);
        quit(1);
    }

    /* This gives the same result as "(kdbx) p processor_ptr" */
    if(!read_sym_addr("processor_ptr", &headptr, &error)){
        fprintf(stderr, "%s\n", error);
        quit(1);
    }

    /* get ncpus */
    if(!read_sym_val("ncpus", NUMBER, &ncpus, &error)){
        fprintf(stderr, "Couldn't read ncpus:\n");
        fprintf(stderr, "%s\n", error);
        quit(1);
    }

    for (i=0; i < ncpus; i++) {

        /* if user wants only one cpu and this is not the one, skip */
        if (cpuflag)
            if (cpuarg != i) continue;

        /* get the ith pointer (values) in the array */
```

Example 3–2: Extension That Uses Linked Lists: callout.c (cont.)

```
    sprintf(buf, "set $hexints=0");
    dbx(buf, False);
    sprintf(buf, "p \*(long \*)0x%lx", headptr+8*i);
    dbx(buf, True);
    if((resp = read_response(&status)) == NULL){
        print_status("Couldn't read value of processor_ptr[i]:", &status);
        quit(1);
    }
    ptr_val = strtoul(resp, (char**)NULL, 10);
    free(resp);

    if (! ptr_val) continue; /* continue if this slot is disabled */

    if(!cast(ptr_val, "struct processor", &processor, &error)){
        fprintf(stderr, "Couldn't cast to a processor:\n");
        fprintf(stderr, "%s:\n", error);
        quit(1);
    }

    if(!read_field_vals(processor, processor_fields, NUM_PROCESSOR_FIELDS)){
        field_errors(processor_fields, NUM_PROCESSOR_FIELDS);
        quit(1);
    }

    if (processor_fields[5].data == 0) continue;

    print("");
    sprintf(buf, "Processor:                %10u", i);
    print(buf);
    sprintf(buf, "Current time (in ticks):        %10u",
processor_fields[4].data ); /*lbolt*/
    print(buf);

    /* for first element, we are interested in time only */

    print("");

    sprintf(buf, "          FUNCTION                ARGUMENT    TICKS(delta)");
    print(buf);
    print( "          =====                =====    =====");

    /* walk through the rest of the list */
    next = (long) processor_fields[3].data;
    while(next != 0) {
        if(!cast(next, "struct callout", &callout, &error)){
            fprintf(stderr, "Couldn't cast to a callout:\n");
            fprintf(stderr, "%s:\n", error);
        }
        if(!read_field_vals(callout, callout_fields, NUM_CALLOUT_FIELDS)){
            field_errors(callout_fields, NUM_CALLOUT_FIELDS);
            break;
        }
        func = addr_to_proc((long) callout_fields[2].data);
        format_addr((long) callout_fields[1].data, arg);
        sprintf(buf, "%-32.32s %12s %12d", func, arg,
((long)callout_fields[0].data & CALLTODO_TIME) -
(long)processor_fields[4].data);
        print(buf);
        next = (long) callout_fields[3].data;
    }
}
```


Example 3–2: Extension That Uses Linked Lists: callout.c (cont.)

```
    }

    } /* end of for */

    quit(0);

} /* end of main() */
```

Example 3–3: Template Extensions Using Arrays

```
#include <stdio.h>
#include <kdbx.h>

static char *help_string =
"<Usage info>                               \\n\ [1]
";

FieldRec fields[] = {
  <data fields> [2]
};
#define NUM_FIELDS (sizeof(fields)/sizeof(fields[0]))
main(argc, argv)
int argc;
char **argv;
{
  int i, size;
  char *error, *ptr;
  DataStruct head, ele;
  check_args(argc, argv, help_string);

  if(!check_fields("<array element type>", fields, NUM_FIELDS, NULL)) { [3]
    field_errors(fields, NUM_FIELDS);
    quit(1);
  }

  if(argc == 0) head = read_sym("<file>"); [4]

  if(!read_sym_val("<symbol containing size of array>", NUMBER, [5]
    (caddr_t *) &size, &error) ||
    fprintf(stderr, "Couldn't read size:\n");
    fprintf(stderr, "%s\n", error);
    quit(1);
  }
  <print header> [6]
  if(argc == 0){
    for(i=0;i<size;i++){
      if((ele = array_element(head, i, &error)) == NULL){
        fprintf(stderr, "Couldn't get array element\n");
        fprintf(stderr, "%s\n", error);
        return(False);
      }
      <print fields in this element> [7]
    }
  }
}
```

Example 3–3: Template Extensions Using Arrays (cont.)

```
}  
}
```

- ❶ The help string is output by the `check_args` function if the user enters the `help extension_name` command at the `kdbx` prompt. The first line of the help string should be a one-line description of the extension. The rest should be a complete description of the arguments. Also, each line should end with the string `\\\n\`.
- ❷ Every structure field to be extracted needs an entry. The first field is the name of the next extracted field; the second field is the type. The last two fields are for output and initialize to `NULL`.
- ❸ Specifies the type of the element in the array.
- ❹ Specifies the variable containing the beginning address of the array.
- ❺ Specifies the variable containing the size of the array. Note that reading variables is only one way to access this information. Other methods include the following:
 - Defining the array size with a `#define` macro call. If you use this method, you need to include the appropriate header file and use the macro in the extension.
 - Querying `dbx` for the array size as follows:

```
dbx("print sizeof(array//sizeof(array[0] ")
```
 - Hard coding the array size.
- ❻ Specifies the string to be displayed as the table header.
- ❼ Extracts, formats, and prints the field information.

Example 3–4: Extension That Uses Arrays: `file.c`

```
#include <stdio.h>  
#include <sys/fcntl.h>  
#include <kdbx.h>  
#include <nlist.h>  
#define SHOW_UTT  
#include <sys/user.h>  
#define KERNEL_FILE  
#include <sys/file.h>  
#include <sys/proc.h>  
  
static char *help_string =  
"file - print out the file table           \\n\  
  Usage : file [addresses...]           \\n\  
  If no arguments are present, all file entries with non-zero reference \\n\  
  counts are printed. Otherwise, the file entries named by the addresses\\n\  
  are printed.                           \\n\  
";
```

Example 3–4: Extension That Uses Arrays: file.c (cont.)

```
char buffer[256];

/* *** Implement addresses *** */

FieldRec fields[] = {
    { ".f_type", NUMBER, NULL, NULL },
    { ".f_count", NUMBER, NULL, NULL },
    { ".f_msgcount", NUMBER, NULL, NULL },
    { ".f_cred", NUMBER, NULL, NULL },
    { ".f_data", NUMBER, NULL, NULL },
    { ".f_ops", NUMBER, NULL, NULL },
    { ".f_u.fu_offset", NUMBER, NULL, NULL },
    { ".f_flag", NUMBER, NULL, NULL }
};

FieldRec fields_pid[] = {
    { ".pe_pid", NUMBER, NULL, NULL },
    { ".pe_proc", NUMBER, NULL, NULL }
};

FieldRec utask_fields[] = {
    { ".uu_file_state.uf_lastfile", NUMBER, NULL, NULL }, /* 0 */
    { ".uu_file_state.uf_ofile", ARRAY, NULL, NULL }, /* 1 */
    { ".uu_file_state.uf_pofile", ARRAY, NULL, NULL }, /* 2 */
    { ".uu_file_state.uf_ofile_of", NUMBER, NULL, NULL }, /* 3 */
    { ".uu_file_state.uf_pofile_of", NUMBER, NULL, NULL }, /* 4 */
    { ".uu_file_state.uf_of_count", NUMBER, NULL, NULL }, /* 5 */
};

#define NUM_FIELDS (sizeof(fields)/sizeof(fields[0]))
#define NUM_UTASK_FIELDS (sizeof(utask_fields)/sizeof(utask_fields[0]))

static char *get_type(int type)
{
    static char buf[5];

    switch(type){
    case 1: return("file");
    case 2: return("sock");
    case 3: return("npip");
    case 4: return("pipe");
    default:
        sprintf(buf, "%3d", type);
        return(buf);
    }
}

long vn_addr, socket_addr;
int proc_size; /* will be obtained from dbx */

static Boolean prfile(DataStruct ele)
{
    char *error, op_buf[12], *ops, buf[256], address[12], cred[12], data[12];

    if(!read_field_vals(ele, fields, NUM_FIELDS)){
        field_errors(fields, NUM_FIELDS);
        return(False);
    }
    if((long) fields[1].data == 0) return(True);
    if((long) (fields[5].data) == 0) ops = " *Null*";
```

Example 3-4: Extension That Uses Arrays: file.c (cont.)

```
else if((long) (fields[5].data) == vn_addr) ops = " vnops";
else if((long) (fields[5].data) == socket_addr) ops = "sockops";
else format_addr((long) fields[5].data, op_buf);
format_addr((long) struct_addr(ele), address);
format_addr((long) fields[3].data, cred);
format_addr((long) fields[4].data, data);
sprintf(buf, "%s %s %4d %4d %s %11s %11s %6d%s%s%s%s%s%s%s",
address, get_type((int) fields[0].data), fields[1].data,
fields[2].data, ops, data, cred, fields[6].data,
(long) fields[7].data) & FREAD ? " r" : "",
(long) fields[7].data) & FWRITE ? " w" : "",
(long) fields[7].data) & FAPPEND ? " a" : "",
(long) fields[7].data) & FNDELAY ? " nd" : "",
(long) fields[7].data) & FMARK ? " m" : "",
(long) fields[7].data) & FDEFER ? " d" : "",
(long) fields[7].data) & FASYNC ? " as" : "",
(long) fields[7].data) & FSHLOCK ? " sh" : "",
(long) fields[7].data) & FEXLOCK ? " ex" : "");
print(buf);
return(True);
}

static Boolean prfiles(DataStruct fil, int n)
{
    DataStruct ele;
    char *error;

    if((ele = array_element(fil, n, &error)) == NULL){
        fprintf(stderr, "Couldn't get array element\n");
        fprintf(stderr, "%s\n", error);
        return(False);
    }
    return(prfile(ele));
}

static void Usage(void){
    fprintf(stderr, "Usage : file [addresses...]\n");
    quit(1);
}

main(int argc, char **argv)
{
    int i;
    long nfile, addr;
    char *error, *ptr, *resp;
    DataStruct fil;
    Status status;

    check_args(argc, argv, help_string);
    argv++;
    argc--;

    if(!check_fields("struct file", fields, NUM_FIELDS, NULL)){
        field_errors(fields, NUM_FIELDS);
        quit(1);
    }
    if(!check_fields("struct pid_entry", fields_pid, 2, NULL){
        field_errors(fields, 2);
        quit(1);
    }
    if(!check_fields("struct utask", utask_fields, NUM_UTASK_FIELDS, NULL){
```

Example 3–4: Extension That Uses Arrays: file.c (cont.)

```
field_errors(fields, NUM_UTASK_FIELDS);
quit(1);
}

if(!read_sym_addr("vnops", &vn_addr, &error) ||
!read_sym_addr("socketops", &socket_addr, &error)){
    fprintf(stderr, "Couldn't read vnops or socketops:\n");
    fprintf(stderr, "%s\n", error);
    quit(1);
}

print("Addr      Type  Ref  Msg Fileops      F_data      Cred Offset
Flags");
print("-----  ----  ---  ---  -----  -----  -----  -----
=====");
if(argc == 0){
    /*
     * New code added to access open files in processes, in
     * the absence of static file table, file, nfile, etc..
     */

    /*
     * get the size of proc structure
     */
    sprintf(buffer, "set $hexints=0");
    dbx(buffer, False);
    sprintf(buffer, "print sizeof(struct proc)");
    dbx(buffer, True);
    if((resp = read_response(&status)) == NULL){
        print_status("Couldn't read sizeof proc", &status);
        proc_size = sizeof(struct proc);
    }
    else
        proc_size = strtoul(resp, (char**)NULL, 10);
    free(resp);

    if ( get_all_open_files_from_active_processes() ) {
        fprintf(stderr, "Couldn't get open files from processes:\n");
        quit(1);
    }
}
else {
    while(*argv){
        addr = strtoul(*argv, &ptr, 16);
        if(*ptr != '\0'){
            fprintf(stderr, "Couldn't parse %s to a number\n", *argv);
            quit(1);
        }
        if(!cast(addr, "struct file", &fil, &error)){
            fprintf(stderr, "Couldn't cast address to a file:\n");
            fprintf(stderr, "%s\n", error);
            quit(1);
        }
        if(!prfile(fil))
            fprintf(stderr, "Continuing with next file address.\n");
        argv++;
    }
}
quit(0);
}
/*
```

Example 3-4: Extension That Uses Arrays: file.c (cont.)

```
* Figure out the location of the utask structure in the supertask
* #define proc_to_utask(p) (long)(p+sizeof(struct proc))
*/

/*
 * Figure out if this a system with the capability of
 * extending the number of open files per process above 64
 */
#ifdef NOFILE_IN_U
# define OFILE_EXTEND
#else
# define NOFILE_IN_U NOFILE
#endif

/*
 * Define a generic NULL pointer
 */
#define NIL_PTR(type) (type *) 0x0

get_all_open_files_from_active_processes()
{
    long pidtab_base;          /* Start address of the process table */
    long npid;                /* Number of processes in the process table */
    char *error;

    if (!read_sym_val("pidtab", NUMBER, &pidtab_base, &error) ||
        !read_sym_val("npid", NUMBER, &npid, &error) ){
        fprintf(stderr, "Couldn't read pid or npid:\n");
        fprintf(stderr, "%s\n", error);
        quit(1);
    }

    if ( check_procs (pidtab_base, npid) )
        return(0);
    else
        return(1);
}

check_procs(pidtab_base, npid)
    long pidtab_base;
    long npid;
{
    int i, index, first_file;
    long addr;
    DataStruct pid_entry_struct, pid_entry_ele, utask_struct, fil;
    DataStruct ofile, pofile;
    char *error;
    long addr_of_proc, start_addr, vall, fp, last_fp;
    char buf[256];

    /*
     * Walk the pid table
     */
    pid_entry_struct = read_sym("pidtab");

    for (index = 0; index < npid; index++)
```

Example 3–4: Extension That Uses Arrays: file.c (cont.)

```
{
    if((pid_entry_ele = array_element(pid_entry_struct, index, &error))==NULL){
        fprintf(stderr, "Couldn't get pid array element %d\n", index);
        fprintf(stderr, "%s\n", error);
        continue;
    }
    if(!read_field_vals(pid_entry_ele, fields_pid, 2)) {
        fprintf(stderr, "Couldn't get values of pid array element %d\n", index);
        field_errors(fields_pid, 2);
        continue;
    }
    addr_of_proc = (long)fields_pid[1].data;
    if (addr_of_proc == 0)
        continue;
    first_file = True;
    addr = addr_of_proc + proc_size;

    if(!cast(addr, "struct utask", &utask_struct, &error)){
        fprintf(stderr, "Couldn't cast address to a utask (bogus?):\n");
        fprintf(stderr, "%s\n", error);
        continue;
    }
    if(!read_field_vals(utask_struct, utask_fields, 3)) {
        fprintf(stderr, "Couldn't read values of utask:\n");
        field_errors(fields_pid, 3);
        continue;
    }
    addr = (long) utask_fields[1].data;
    if (addr == NULL)
        continue;

    for(i=0;i<=(int)utask_fields[0].data;i++){
        if(i>=NOFILE_IN_U){
if (utask_fields[3].data == NULL)
            continue;
        start_addr = (long)((long *)utask_fields[3].data + i-NOFILE_IN_U) ;
        if(!read_memory(start_addr , sizeof(struct file *), (char *)&vall,
&error)) {
            fprintf(stderr, "Start addr:0x%lx bytes:%d\n", start_addr, sizeof(long
*));
            fprintf(stderr, "Couldn't read memory for extn files: %s\n", error);
            continue;
        }
        }
        else {
ofile = (DataStruct) utask_fields[1].data;
pofile = (DataStruct) utask_fields[2].data;
        }
        if (i < NOFILE_IN_U)
if(!array_element_val(ofile, i, &vall, &error)){
            fprintf(stderr, "Couldn't read %d'th element of ofile|pofile:\n", i);
            fprintf(stderr, "%s\n", error);
            continue;
        }
        }
        fp = vall;
        if(fp == 0) continue;
        if(fp == last_fp) continue; /* eliminate duplicates */
        last_fp = fp;
        if(!cast(fp, "struct file", &fil, &error)){
            fprintf(stderr, "Couldn't cast address to a file:\n");
            fprintf(stderr, "%s\n", error);
        }
    }
}
```

Example 3–4: Extension That Uses Arrays: file.c (cont.)

```
quit(1);
    }
    if (first_file) {
sprintf(buf, "[Process ID: %d]", fields_pid[0].data);
print(buf);
first_file = False;
    }
    if(!prfile(fil))
fprintf(stderr, "Continuing with next file address.\n");
    }
} /* for loop */

return(True);
} /* end */
```

Example 3–5: Extension That Uses Global Symbols: sum.c

```
#include <stdio.h>
#include <kdbx.h>

static char *help_string =
"sum - print a summary of the system          \\n\\n
  Usage : sum                                \\n\\n";

static void read_var(name, type, val)
char *name;
int type;
long *val;
{
    char *error;
    long n;

    if(!read_sym_val(name, type, &n, &error)){
        fprintf(stderr, "Reading %s:\n", name);
        fprintf(stderr, "%s\n", error);
        quit(1);
    }
    *val = n;
}

main(argc, argv)
int argc;
char **argv;
{
    DataStruct utsname, cpup, time;
    char buf[256], *error, *resp, *sysname, *release, *version, *machine;
    long avail, secs, tmp;

    check_args(argc, argv, help_string);
    read_var("utsname.nodename", STRING, &resp);
    sprintf(buf, "Hostname : %s", resp);
    print(buf);
    free(resp);
    read_var("ncpus", NUMBER, &avail);
} /*
```


Example 3–5: Extension That Uses Global Symbols: `sum.c` (cont.)

```
* cpup no longer exists, emulate platform_string(),
* a.k.a. get_system_type_string().
read_var("cpup.system_string", STRING, &resp);
*/
read_var("rpb->rpb_vers", NUMBER, &tmp);
if (tmp < 5)
    resp = "Unknown System Type";
else
    read_var(
(char *)rpb + rpb->rpb_dsr_off + "
((struct rpb_dsr *)"
 ((char *)rpb + rpb->rpb_dsr_off)->rpb_sysname_off + sizeof(long)",
    STRING, &resp);
sprintf(buf, "cpu: %s\tavail: %d", resp, avail);
print(buf);
free(resp);
read_var("boottime.tv_sec", NUMBER, &secs);
sprintf(buf, "Boot-time:\t%s", ctime(&secs));
buf[strlen(buf) - 1] = '\0';
print(buf);
read_var("time.tv_sec", NUMBER, &secs);
sprintf(buf, "Time:\t%s", ctime(&secs));
buf[strlen(buf) - 1] = '\0';
print(buf);
read_var("utsname.sysname", STRING, &sysname);
read_var("utsname.release", STRING, &release);
read_var("utsname.version", STRING, &version);
read_var("utsname.machine", STRING, &machine);
sprintf(buf, "Kernel : %s release %s version %s (%s)", sysname, release,
    version, machine);
print(buf);
quit(0);
}
```

3.4 Compiling Custom Extensions

After you have written the extension, you need to compile it. To compile the extension, enter the following command:

```
% cc -o test test.c -lkdbx
```

This `cc` command compiles an extension named `test.c`. The `kdbx.a` library is linked with the extensions, as specified by the `-l` flag. The output from this command is named `test`, as specified by the `-o` flag.

Once the extension compiles successfully, you should test it and, if necessary, debug it as described in Section 3.5.

When the extension is ready for use, place it in a directory that is accessible to other users. Extensions provided with the operating system are located in the `/var/kdbx` directory.

The following example shows how to invoke the `test` extension from within the `kdbx` debugger:

```
# kdbx -k /vmunix
dbx version 5.0
Type 'help' for help.

(kdbx) test
Hostname : system.dec.com
cpu: DEC3000 - M500      avail: 1
Boot-time:      Fri Nov  6 16:09:10 1992
Time:   Mon Nov  9 10:51:48 1992
Kernel : OSF1 release 1.2 version 1.2 (alpha)
(kdbx)
```

3.5 Debugging Custom Extensions

The `kdbx` debugger and the `dbx` debugger include the capability to communicate with each other using two named pipes. The task of debugging an extension is easier if you use a workstation with two windows or two terminals. In this way, you can dedicate one window or terminal to the `kdbx` debugger and one window or terminal to the `dbx` debugger. However, you can debug an extension from a single terminal.

This section explains how to begin your `kdbx` and `dbx` sessions when you have two windows or terminals and when you have a single terminal. The examples illustrate debugging the `test` extension that was compiled in Section 3.4.

If you are using a workstation with two windows or have two terminals, perform the following steps to set up your `kdbx` and `dbx` debugging sessions:

1. Open two sessions: one running `kdbx` on the running kernel and the other running `dbx` on the source file for the custom extension `test` as follows:

Begin the `kdbx` session:

```
# kdbx -k /vmunix
dbx version 5.0
Type 'help' for help.

stopped at [thread_block:1440 ,0xfffffc00002de5b0]  Source not available
```

Begin the `dbx` session:

```
# dbx test
dbx version 5.0
Type 'help' for help.

(dbx)
```

2. Set up `kdbx` and `dbx` to communicate with each other. In the `kdbx` session, enter the `procpd` alias to create the files `/tmp/pipein` and `/tmp/pipeout` as follows:

```
(kdbx) procpd
```

The file `pipein` directs output from the `dbx` session to the `kdbx` session. The file `pipeout` directs output from the `kdbx` session to the `dbx` session.

3. In the `dbx` session, enter the `run` command to execute the `test` extension in the `kdbx` session, specifying the files `/tmp/pipein` and `/tmp/pipeout` on the command line as follows:

```
(dbx) run < /tmp/pipeout > /tmp/pipein
```

4. As you step through the extension in the `dbx` session, you see the results of any action in the `kdbx` session. At this point, you can use the available `dbx` commands and options.

If you are using one terminal, perform the following steps to set up your `kdbx` and `dbx` sessions:

1. Issue the following command to invoke `kdbx` with the debugging environment:

```
# echo 'procpd' | kdbx -k /vmunix &
dbx version 5.0
Type 'help' for help.

stopped at [thread_block:1403 ,0xfffffc000032d860]   Source not available
#
```

2. Invoke the `dbx` debugger as follows:

```
# dbx test
dbx version 5.0
Type 'help' for help.

(dbx)
```

3. As you step through the extension in the `dbx` session, you see the results of any action in the `kdbx` session. At this point, you can use the available `dbx` commands and options.

Crash Analysis Examples

Finding problems in crash dump files is a task that takes practice and experience to do well. Exactly how you determine what caused a crash varies depending on how the system crashed. The cause of some crashes is relatively easy to determine, while finding the cause of other crashes is difficult and time-consuming.

This chapter helps you analyze crash dump files by providing the following information:

- Guidelines for examining crash dump files (Section 4.1)
- Examples of identifying the cause of a software panic (Section 4.2)
- Examples of identifying the cause of a hardware trap (Section 4.3)
- An example of finding a panic string that is not in the current thread (Section 4.4)
- An example of identifying the cause of a crash on an SMP system (Section 4.5)

For information about how crash dump files are created, see the *System Administration* manual.

4.1 Guidelines for Examining Crash Dump Files

In examining crash dump files, there is no one way to determine the cause of a system crash. However, following these steps should help you identify the events that lead to most crashes:

1. Gather some facts about the system; for example, operating system type, version number, revision level, hardware configuration.
2. Locate the thread executing at the time of the crash. Most likely, this thread contains the events that lead to the panic.
3. Look at the panic string, if one exists. This string is contained in the preserved message buffer (`pmsgbuf`) and in the `panicstr` global variable. The panic string gives a reason for the crash.
4. Identify the function that called the `panic` or `trap` function. That function is the one that caused the system to crash.
5. Examine the source code for the function that caused the crash to infer the error that caused the crash. You might also need to examine

related data structures and functions that appear earlier in the stack. An earlier function might have passed corrupt data to the function that caused a crash.

6. Determine whether you can fix the problem.

If the system crashed because of a hardware problem (for example, because a memory board became corrupt), correcting the problem probably requires repairing or replacing the hardware. You might be able to disconnect the hardware that caused the problem and operate without it until it is repaired or replaced. If you need to repair or replace hardware, call your support representative.

If a software panic caused the crash, you can fix the problem if it is in software you or someone else at your company wrote. Otherwise, you must request that the producer of the software fix the problem by calling your support representative.

4.2 Identifying a Crash Caused by a Software Problem

When software encounters a state from which it cannot continue, it calls the system `panic` function. For example, if the software attempts to access an area of memory that is protected from access, the software might call the `panic` function and crash the system.

In most cases, only system programmers can fix the problem that caused a panic because most panics are caused by software errors. However, some system panics reflect other problems. For example, if a memory board becomes corrupted, software that attempts to write to that board might call the `panic` function and crash the system. In this case, the solution might be to replace the memory board and reboot the system.

The sections that follow demonstrate finding the cause of a software panic using the `dbx` and `kdbx` debuggers. You can also examine output from the `crashdc` crash data collection tool to help you determine the cause of a crash. Sample output from `crashdc` is shown and explained in Appendix A.

4.2.1 Using `dbx` to Determine the Cause of a Software Panic

The following example shows a method for identifying a software panic with the `dbx` debugger:

```
# dbx -k vmunix.0 vmzcore.0
dbx version 5.0
Type 'help' for help.

stopped at [boot:753 ,0xfffffc00003c4ae4] Source not available

(dbx) p panicstr [1]
0xfffffc000044b648 = "ialloc: dup alloc"
(dbx) t [2]
> 0 boot(paniced = 0, arghowto = 0) ["../../../../src/kernel/arch/alpha/machdep.\
```

```

c":753, 0xfffffc00003c4ae4]
  1 panic(s = 0xfffffc000044b618 = "mode = 0%, inum = %d, pref = %d fs = %s\n")\
  ["../../../../src/kernel/bsd/subr_prf.c":1119, 0xfffffc00002bdbb0]
  2 ialloc(pip = 0xffffffff8c6acc40, ipref = 57664, mode = 0, ipp = 0xffffffff8c\
  f95af8) ["../../../../src/kernel/ufs/ufs_alloc.c":501, 0xfffffc00002dab48]
  3 maknode(vap = 0xffffffff8cf95c50, ndp = 0xffffffff8cf922f8, ipp = 0xffffffff\
  8cf95b60) ["../../../../src/kernel/ufs/ufs_vnops.c":2842, 0xfffffc00002ea500]
  4 ufs_create(ndp = 0xffffffff8cf922f8, vap = 0xfffffc00002fe0a0) ["../../../../\
  /src/kernel/ufs/ufs_vnops.c":602, 0xfffffc00002e771c]
  5 vn_open(ndp = 0xffffffff8cf95d18, fmode = 4618, cmode = 416) ["../../../../s\
  rc/kernel/vfs/vfs_vnops.c":258, 0xfffffc00002fel38]
  6 copen(p = 0xffffffff8c6efba0, args = 0xffffffff8cf95e50, retval = 0xffffffff\
  8cf95e40, compat = 0) ["../../../../src/kernel/vfs/vfs_syscalls.c":1379, 0xfffffc\
  00002fb890]
  7 open(p = 0xffffffff8cf95e40, args = (nil), retval = 0x7f4) ["../../../../src\
  /kernel/vfs/vfs_syscalls.c":1340, 0xfffffc00002fb7bc]
  8 syscall(ep = 0xffffffff8cf95ef8, code = 45) ["../../../../src/kernel/arch/al\
  pha/syscall_trap.c":532, 0xfffffc00003cfa34]
  9 _Xsyscall() ["../../../../src/kernel/arch/alpha/locore.s":703, 0xfffffc00003\
  c31e0]
(dbx) q

```

- ❶ Display the panic string (`panicstr`). The panic string shows that the `ialloc` function called the panic function.
- ❷ Perform a stack trace. This confirms that the `ialloc` function at line 501 in file `ufs_alloc.c` called the panic function.

4.2.2 Using `kdbx` to Determine the Cause of a Software Panic

The following example shows a method of finding a software panic with the `kdbx` debugger:

```

# kdbx -k vmunix.3 vmzcore.3
dbx version 5.0
Type 'help' for help.

stopped at [boot:753,0xfffffc00003c4b04] Source not available

(kdbx) sum ❶
Hostname : system.dec.com
cpu: Digital AlphaStation 600 5/266 avail: 1
Boot-time: Tue Oct 6 15:16:41 1998
Time: Tue Oct 27 13:52:11 1998
Kernel : OSF1 release V5.0 version 688.2 (alpha)
(kdbx) p panicstr❷
0xfffffc0000453ea0 = "wdir: compact2"
(kdbx) t ❸
> 0 boot(panicd = 0, arghowto = 0) ["../../../../src/kernel/arch/alpha/machdep\
.c":753, 0xfffffc00003c4b04]
  1 panic(s = 0xfffffc00002e0938 = "p") ["../../../../src/kernel/bsd/subr_prf.c"\
:1119, 0xfffffc00002bdbb0]
  2 direnter(ip = 0xffffffff00000000, ndp = 0xffffffff9d38db60) ["../../../../sr\
c/kernel/ufs/ufs_lookup.c":986, 0xfffffc00002e2adc]
  3 ufs_mkdir(ndp = 0xffffffff9d38a2f8, vap = 0x100000020) ["../../../../src/ker\
nel/ufs/ufs_vnops.c":2383, 0xfffffc00002e9cbc]
  4 mkdir(p = 0xffffffff9c43d7c0, args = 0xffffffff9d38de50, retval = 0xffffffff\
9d38de40) ["../../../../src/kernel/vfs/vfs_syscalls.c":2579, 0xfffffc00002fd930]
  5 syscall(ep = 0xffffffff9d38def8, code = 136) ["../../../../src/kernel/arch/a\
lpha/syscall_trap.c":532, 0xfffffc00003cfa54]
  6 _Xsyscall() ["../../../../src/kernel/arch/alpha/locore.s":703, 0xfffffc00003\

```

```
c3200]
(kdbx) q
dbx (pid 29939) died. Exiting...
```

- 1 Use the `sum` command to get a summary of the system.
- 2 Display the panic string (`panicstr`).
- 3 Perform a stack trace of the current thread block. The stack trace shows that the `direnter` function, at line 986 in file `ufs_lookup.c`, called the `panic` function.

4.3 Identifying a Hardware Exception

Occasionally, your system might crash due to a hardware error. During a hardware exception, the hardware encounters a situation from which it cannot continue. For example, the hardware might detect a parity error in a portion of memory that is necessary for its successful operation. When a hardware exception occurs, the hardware stores information in registers and stops operation. When control returns to the software, it normally calls the `panic` function and the system crashes.

The sections that follow show how to identify hardware traps using the `dbx` and `kdbx` debuggers. You can also examine output from the `crashdc` crash data collection tool to help you determine the cause of a crash. Sample output from `crashdc` is shown and explained in Appendix A.

4.3.1 Using `dbx` to Determine the Cause of a Hardware Error

The following example shows a method for identifying a hardware trap with the `dbx` debugger:

```
# dbx -k vmunix.1 vmzcore.1
dbx version 5.0
Type 'help' for help.
(dbx) sh strings vmunix.1 | grep '(Rev' 1
Tru64 UNIX V5.0-1 (Rev. 961); Wed Mar 18 16:12:36 EST 1999

(dbx) p utsname 2
struct {
    sysname = "OSF1"
    nodename = "system.dec.com"
    release = "V5.0"
    version = "961"
    machine = "alpha"
}

(dbx) p panicstr 3
0xfffffc0000489350 = "trap: Kernel mode prot fault\n"

(dbx) t 4
> 0 boot(paniced = 0, arghowto = 0) ["/usr/sde/alpha/build/alpha.nightly/src/ker\
nel/arch/alpha/machdep.c":
    1 panic(s = 0xfffffc0000489350 = "trap: Kernel mode prot fault\n") ["/usr/sde\
/alpha/build/alpha.nightly/src/kernel/bsd/subr_prf.c":1099, 0xfffffc00002c0730]
    2 trap() ["/usr/sde/alpha/build/alpha.nightly/src/kernel/arch/alpha/trap.c":54\
```



```
4, 0xfffffc00003e0c78]
  3 _XentMM() ["/usr/sde/alpha/build/alpha.nightly/src/kernel/arch/alpha/locore.\
s":702, 0xfffffc00003d4ff4]
```

```
(dbx) kps 5
```

```
  PID  COMM
00000  kernel idle
00001  init
00002  device server
00003  exception hdlr
00663  ypbind
00018  cfgmgr
00219  automount
:
:
00265  cron
00293  xdm
02311  inetd
00278  lpd
01443  csh
01442  rlogind
01646  rlogind
01647  csh
```

```
(dbx) p $pid 6
```

```
2311
```

```
(dbx) p *pmsgbuf 7
```

```
struct {
  msg_magic = 405601
  msg_bufx = 62
  msg_bufc = 3825
  msg_bufc = "unknown flag
printstate: unknown flag
printstate: unknown flag
de: table is full
<3>vnode: table is full
:
:
```

```
<3>arp: local IP address 0xffffffff82b40429 in use by
hardware address 08:00:2B:20:19:CD
<3>arp: local IP address 0xffffffff82b40429 in use by
hardware address 08:00:2B:2B:F6:3B
```

```
va=0000000000000028, status word=0000000000000000, pc=fffffc000032972c
panic: trap: Kernel mode prot fault
syncing disks... 3 3 done
printstate: unknown flag
printstate: unknown flag
printstate: unknown flag
printstate: unknown flag
printstate: u"
}
```

```
(dbx) px savedefp
0xffffffff89b2b4e0
```

```
(dbx) p savedefp
0xffffffff89b2b4e0
```

```
(dbx) p savedefp[28]
18446739675666356012
```

```

(dbx) px savedefp[28] 8
0xffffffff000032972c

(dbx) savedefp[28]/i 9
[nfs_putpage:2344, 0xffffffff000032972c]      ld1      r5, 40(r1)
(dbx) savedefp[23]/i 10
[ubc_invalidate:1768, 0xffffffff0000315fe0]    st1      r0, 84(sp)

(dbx) func nfs_putpage 11
(dbx) file 12
/usr/sde/alpha/build/alpha.nightly/src/kernel/kern/sched_prim.c
(dbx) func ubc_invalidate 13
ubc_invalidate: Source not available

(dbx) file 14
/usr/sde/alpha/build/alpha.nightly/src/kernel/vfs/vfs_ubc.c

(dbx) q

```

- 1 You can use the `sh` command to enter commands to the shell. In this case, enter the `strings` and `grep` commands to pull the operating system revision number in the `vmunix.1` dump file.
- 2 Display the `utsname` structure to obtain more information about the operating system version.
- 3 Display the panic string (`panicstr`). The panic function was called by a trap function.
- 4 Perform a stack trace. This confirms that the trap function called the panic function. However, the stack trace does not show what caused the trap.
- 5 Look to see what processes were running when the system crashed by entering the `kps` command.
- 6 Look to see what the process ID (PID) was pointing to at the time of the crash. In this case, the PID was pointing to process 2311, which is the `inetd` daemon, from the `kps` command output.
- 7 Display the preserved message buffer (`pmsgbuf`). Note that this buffer contains the program counter (`pc`) value, which is displayed in the following line:

```
va=0000000000000028, status word=0000000000000000, pc=ffffffc000032972c
```
- 8 Display register 28 of the exception frame pointer (`savedefp`). This register always contains the `pc` value. You can always obtain the `pc` value from either the preserved message buffer or register 28 of the exception frame pointer.
- 9 Disassemble the `pc` to determine its contents. The `pc` at the time of the crash contained the `nfs_putpage` function at line 2344.
- 10 Disassemble the return address to determine its contents. The return value at the time of the crash contained the `ubc_invalidate` function at line 1768.

- 11 Point the dbx debugger to the `nfs_putpage` function.
- 12 Display the name of the source file that contains the `nfs_putpage` function.
- 13 Point the dbx debugger to the `ubc_invalidate` function.
- 14 Display the name of the source file that contains the `ubc_invalidate` function.

The result from this example shows that the `ubc_invalidate` function, which resides in the `/vfs/vfs_ubc.c` file at line number 1768, called the `nfs_putpage` function at line number 2344 in the `/kern/sched_prim.c` file and the system stopped.

4.3.2 Using `kdbx` to Determine the Cause of a Hardware Error

The following example shows a method for identifying a hardware error with the `kdbx` debugger:

```
# kdbx -k vmunix.5 vmzcore.5
dbx version 5.0
Type 'help' for help.

stopped at [boot:753 ,0xffffffff00003c4b04] Source not available
(kdbx) sum 1
Hostname : system.dec.com
cpu: Digital AlphaStation 600 5/266      avail: 1
Boot-time:      Tue Oct  6 15:16:41 1998
Time:      Tue Oct 27 13:52:11 1998
Kernel : OSF1 release V5.0 version 688.2 (alpha)
(kdbx) p panicstr 2
0xffffffff0000471030 = "ECC Error"
(kdbx) t 3
> 0 boot(paniced = 0, arghowto = 0) ["../../../../src/kernel/arch/alpha/machdep.\
c":753, 0xffffffff00003c4b04]
  1 panic(s = 0x670) ["../../../../src/kernel/bsd/subr_prf.c":1119, 0xffffffff00002\
bdbb0]
  2 kn15aa_machcheck(type = 1648, cmcf = 0xffffffff0000f8050 = , framep = 0xffff\
ffff94f79ef8) ["../../../../src/kernel/arch/alpha/hal/kn15aa.c":1269, 0xffffffff000\
03da62c]
  3 mach_error(type = -1795711240, phys_logout = 0x3, regs = 0x6) ["../../../../s\
rc/kernel/arch/alpha/hal/cpusw.c":323, 0xffffffff00003d7dc0]
  4 _XentInt() ["../../../../src/kernel/arch/alpha/locore.s":609, 0xffffffff00003c3\
148]
(kdbx) q
dbx (pid 337) died.  Exiting...
```

- 1 Use the `sum` command to get a summary of the system.
- 2 Display the panic string (`panicstr`).
- 3 Perform a stack trace. Because the `kn15aa_machcheck` function (which is a hardware checking function) called the `panic` function, the system crash was probably the result of a hardware error.

4.4 Finding a Panic String in a Thread Other Than the Current Thread

The dbx and kdbx debuggers have the concept of the current thread. In many cases, when you invoke one of the debuggers to analyze a crash dump, the panic string is in the current thread. At times, however, the current thread contains no panic string and so is probably not the thread that caused the crash.

The following example shows a method for stepping through kernel threads to identify the events that lead to the crash:

```
# dbx -k ./vmunix.2 ./vmzcore.2
dbx version 5.0
Type 'help' for help.
thread 0x8d431c68 stopped at [thread_block:1305 +0x114,0xfffffc000033961c] \
  Source not available
(dbx) p panicstr 1
0xfffffc000048a0c8 = "kernel memory fault"
(dbx) t 2

> 0 thread_block() ["../../../../src/kernel/kern/sched_prim.c":1305, 0xfffffc0\
e
00033961c]
  1 mpsleep(chan = 0xfffffffff8d4ef450 = , pri = 282, wmesg = 0xfffffc000046f\
290 = "network", timo = 0, lockp = (nil), flags = 0) ["../../../../src/kernel/\
bsd/kern_synch.c":267, 0xfffffc00002b772c]
  2 sosleep(so = 0xfffffffff8d4ef408, addr = 0xfffffffff906cfcf4 = "^P", pri = 2 \
82,tmo = 0) ["../../../../src/kernel/bsd/uipc_socket2.c":612, 0xfffffc00002d3784]
  3 accept1(p = 0xfffffffff8f8bfde8, args = 0xfffffffff906cfe50, retval = 0xffff \
ffff906cfe40, compat_43 = 1) ["../../../../src/kernel/bsd/uipc_syscalls.c":300 \
, 0xfffffc00002d4c74]
  4 oaccept(p = 0xfffffffff8d431c68, args = 0xfffffffff906cfe50, retval = 0xffff \
ffff906cfe40) ["../../../../src/kernel/bsd/uipc_syscalls.c":250, 0xfffffc00002d\
4b0c]
  5 syscall(ep = 0xfffffffff906cfef8, code = 99, sr = 1) ["../../../../src/kern \
el/arch/alpha/syscall_trap.c":499, 0xfffffc00003ec18c]
  6 _Xsyscall() ["../../../../src/kernel/arch/alpha/locore.s":675, 0xfffffc000\
03df96c]
(dbx) tlist 3
thread 0x8d431a60 stopped at [thread_block:1305 +0x114,0xfffffc000033961c] \
Source not available
thread 0x8d431858 stopped at [thread_block:1289 +0x18,0xfffffc00003394b8] \
Source not available
thread 0x8d431650 stopped at [thread_block:1289 +0x18,0xfffffc00003394b8] \
Source not available
thread 0x8d431448 stopped at [thread_block:1305 +0x114,0xfffffc000033961c] \
Source not available
thread 0x8d431240 stopped at [thread_block:1305 +0x114,0xfffffc000033961c] \
Source not available
:
:
thread 0x8d42f5d0 stopped at [boot:696 ,0xfffffc00003e119c] Source not
\
available
thread 0x8d42f3c8 stopped at [thread_block:1289 +0x18,0xfffffc00003394b8] \
Source not available
thread 0x8d42f1c0 stopped at [thread_block:1289 +0x18,0xfffffc00003394b8] \
Source not available
thread 0x8d42efb8 stopped at [thread_block:1289 +0x18,0xfffffc00003394b8] \
```

```

Source not available
thread 0x8d42dd70 stopped at [thread_block:1289 +0x18,0xfffffc00003394b8] \
Source not available
(dbx) tset 0x8d42f5d0 [4]
thread 0x8d42f5d0 stopped at [boot:696 ,0xfffffc00003e119c] Source not ava\
ilable
(dbx) t [5]
> 0 boot(paniced = 0, arghowto = 0) ["/usr/src/kernel/arch/alpha/mac\
hdep.c":694, 0xfffffc00003e1198]
  1 panic(s = 0xfffffc000048a098 = "    sp contents at time of fault: 0x%l01\
6x\r\n\n") ["/usr/src/kernel/bsd/subr_prf.c":1110, 0xfffffc00002beef4]
  2 trap() ["/usr/src/kernel/arch/alpha/trap.c":677, 0xfffffc00003ecc70]
  3 _XentMM() ["/usr/src/kernel/arch/alpha/locore.s":828, 0xfffffc000\
03dfb1c]
  4 pmap_release_page(pa = 18446744071785586688) ["/usr/src/kernel/ar\
ch/alpha/pmap.c":640, 0xfffffc00003e3ecc]
  5 put_free_ptepage(page = 5033216) ["/usr/src/kernel/arch/alpha/pma\
p.c":534, 0xfffffc00003e3ca0]
  6 pmap_destroy(map = 0xfffffff8d5bc428) ["/usr/src/kernel/arch/alp\
ha/pmap.c":1891, 0xfffffc00003e6140]
  7 vm_map_deallocate(map = 0xfffffff81930ee0) ["/usr/src/kernel/vm/\
vm_map.c":482, 0xfffffc00003d03c0]
  8 task_deallocate(task = 0xfffffff8d568d48) ["/usr/src/kernel/kern\
/task.c":237, 0xfffffc000033c1dc]
  9 thread_deallocate(thread = 0x4e4360) ["/usr/src/kernel/kern/threa\
d.c":689, 0xfffffc000033d83c]
 10 reaper_thread() ["/usr/src/kernel/kern/thread.c":1952, 0xfffffc00\
0033e920]
 11 reaper_thread() ["/usr/src/kernel/kern/thread.c":1901, 0xfffffc00\
0033e8ac]
(dbx) q

```

- 1 Display the panic string (`panicstr`) to view the panic message, if any. This message indicates that a memory fault occurred.
- 2 Perform a stack trace of the current thread. Because this thread does not show a call to the `panic` function, you need to look at other threads.
- 3 Examine the system's threads. The thread most likely to contain the panic is the `boot` thread because the `boot` function always executes immediately before the system crashes. If the `boot` thread does not exist, you must examine every thread of every process in the process list.
- 4 Point `dbx` to the `boot` thread at address `0x8d42f5d0`.
- 5 In this example, the problem is in the `pmap_release_page` function at line 640 of the `pmap.c` file.

4.5 Identifying the Cause of a Crash on an SMP System

If you are analyzing crash dump files from an SMP system, you must first determine on which CPU the panic occurred. You can then continue crash dump analysis as you would on a single processor system.

The following example shows a method for determining which CPU caused the crash and which function called the `panic` function:

```

% dbx -k ./vmunix.1 ./vmzcore.1
dbx version 5.0
Type 'help' for help.
stopped at [boot:1494 ,0xffffffff0000442918] Source not available
(dbx) p ustsname [1]
struct {
    sysname = "OSF1"
    nodename = "system.dec.com"
    release = "V5.0"
    version = "688.2"
    machine = "alpha"
}

(dbx) print paniccpu [2]
0
(dbx) p machine_slot[1] [3]
struct {
    is_cpu = 1
    cpu_type = 15
    cpu_subtype = 3
    running = 1
    cpu_ticks = {
        [0] 416162
        [1] 83260
        [2] 1401080
        [3] 11821212
        [4] 1095581
    }
    clock_freq = 1024
    error_restart = 0
    cpu_panicstr = 0xffffffff000059f6a0 = "cpu_ip_intr: panic request"
    cpu_panic_thread = 0xfffffffff8109a780
}

(dbx) p panicstr [4]
0xffffffff0000558ad0 = "simple_lock: uninitialized lock"
(dbx) tset active_threads[paniccpu] [5]
stopped at [boot:1494 ,0xffffffff0000442918]
(dbx) t [6]
> 0 boot(0x0, 0x4, 0xac35c0000000a, 0xffffffff00004403fc, 0xffffffff000000000e) \
["../../../../src/kernel/arch/alpha/machdep.c":1494, 0xffffffff0000442918]
1 panic(s = 0xffffffff0000558b40 = "simple_lock: hierarchy violation") ["..\
2 simple_lock_fault(slp = 0xffffffff00006292f0, state = 0, caller = 0xffffffff\
000046f384, arg = 0xffffffff0000534fd8 = "session.s_fggrp_lock", fmt = 0xffffffff\
0000558de8 = " class already locked: %s\n", error = 0xffffffff0000558b40 = "\
simple_lock: hierarchy violation") ["../../../../src/kernel/kern/lock.c":1558\
, 0xffffffff00003c34ec]
3 simple_lock_hierarchy_violation(slp = 0xffffffff000046f384, state = 184467\
39675668500440, caller = 0xffffffff0000558de8, curhier = 5606208) ["../../../../\
/src/kernel/kern/lock.c":1616, 0xffffffff00003c3620]
4 xnaintr(0xffffffff00005a5158, 0x2, 0xfffffffffb53ef238, 0xffffffff000068a754,\
0xffffffff000055891d) ["../../../../src/kernel/io/dec/netif/if_xna.c":1077, 0x\
ffffff000046f384]
5 _XentInt(0x2, 0xffffffff0000447174, 0xffffffff00005b7d40, 0x2, 0x0) ["../../../../\
6 swap_ipl(0x2, 0xffffffff0000447174, 0xffffffff00005b7d40, 0x2, 0x0) ["../../../../\
7 boot(0x0, 0x0, 0xfffffffffa52c6000, 0xfffffffffb53ef1f8, 0xffffffff00003bf4f\
c) ["../../../../src/kernel/arch/alpha/machdep.c":1434, 0xffffffff000044280c]
8 panic(s = 0xffffffff0000558ad0 = "simple_lock: uninitialized lock") ["..\
9 simple_lock_fault(slp = 0xfffffffffa52c6000, state = 1719, caller = 0xff\
ffc00003734c4, arg = (nil), fmt = (nil), error = 0xffffffff0000558ad0 = "simple\
_lock: uninitialized lock") ["../../../../src/kernel/kern/lock.c":1558, 0xff\
ffc00003c34ec]
10 simple_lock_valid_violation(slp = 0xffffffff00003734c4, state = 0, caller \
= (nil)) ["../../../../src/kernel/kern/lock.c":1584, 0xffffffff00003c3578]

```

```

11 pgrp_ref(0xfffffffffa52c6000, 0x0, 0xfffffc000023ee20, 0x6b7, 0xfffffc000\
05e1080) ["/../../../../src/kernel/bsd/kern_proc.c":561, 0xfffffc00003734c4]
12 exit(0xffffffffffb53ef740, 0x100, 0x1, 0xfffffffffa42e5e80, 0x1) ["/../../../../\
../../../../src/kernel/bsd/kern_exit.c":868, 0xfffffc000023ef30]
13 rexit(0xffffffffff814d2d80, 0xffffffffffb53ef758, 0xffffffffffb53ef8b8, 0x1000\
00001, 0x0) ["/../../../../src/kernel/bsd/kern_exit.c":546, 0xfffffc000023e7dc]
14 syscall(0xffffffffffb53ec000, 0xfffffc000068a300, 0x0, 0x51, 0x1) ["/../../../../\
15 _Xsyscall(0x8, 0x3ff800e6938, 0x14000d0f0, 0x1, 0x11ffffc18) ["/../../../../\
(dbx) p *pmsgbuf 7
struct {
    msg_magic = 405601
    msg_bufl = 701
    msg_bufr = 134
    msg_bufc = "0.64.143, errno 22
NFS server: stale file handle fs(742,645286) file 573 gen 32779
getattr, client address = 16.140.64.143, errno 22

simple_lock: uninitialized lock

pc of caller:      0xfffffc00003734c4
lock address:      0xfffffffffa52c6000
lock class name:   (unknown_simple_lock)
current lock state: 0x00000000e0e9b04a (cpu=0,pc=0xfffffc00e0e9b048,free)

panic (cpu 0): simple_lock: uninitialized lock

simple_lock: hierarchy violation

pc of caller:      0xfffffc000046f384
lock address:      0xfffffc00006292f0
lock info addr:    0xfffffc0000672cc0
lock class name:   xna_softc.lk_xna_softc
class already locked: session.s_fgprp_lock
:
}
(dbx) quit

```

- 1 Display the `ustname` structure to obtain information about the system.
- 2 Display the number of the CPU on which the panic occurred, in this case CPU 0 was the CPU that started the system panic.
- 3 Display the `machine_slot` structure for a CPU other than the one that started the system panic. Notice that the panic string contains:

```
cpu_ip_intro: panic_request
```
- 4 Display the panic string, which in this case indicates that a process attempted to obtain an uninitialized lock.
- 5 Set the context to the CPU that caused the system panic to begin.
- 6 Perform a stack trace on the CPU that started the system panic.

Notice that the `panic` function appears twice in the stack trace. The series of events that resulted in the first call to the `panic` function caused the crash. The events that occurred after the first call to the

`panic` function were performed after the system was corrupt and during an attempt to save data. Normally, any events that occur after the initial call to the `panic` function will not help you determine why the system crashed.

In this example, the problem is in the `pgrp_ref` function on line 561 in the `kern_proc.c` file.

If you follow the stack trace after the `pgrp_ref` function, you can see that the `pgrp_ref` function calls the `simple_lock_valid_violation` function. This function displays information about simple locks, which might be helpful in determining why the system crashed.

- 7 Retrieve the information from the `simple_lock_valid_violation` function by displaying the preserved message buffer.

A

Output from the crashdc Command

This appendix contains a sample `crash-data.n` file created by the `crashdc` command (using a compressed crash-dump file, `vmzcore.0`). The output is explained in the list following the example.

```
#
# Crash Data Collection (Version 1.4)
#
_crash_data_collection_time: Fri Jul 10 01:25:31 EDT 1998 [1]
_current_directory: /
_crash_kernel: /var/adm/crash/vmunix.0
_crash_core: /var/adm/crash/vmzcore.0
_crash_arch: alpha
_crash_os: Tru64 UNIX
_host_version: Tru64 UNIX V5.0 (Rev. 1039); Tue Jun 30 08:26:03 EDT 1998
_crash_version: Tru64 UNIX V5.0 (Rev. 1039); Tue Jun 30 08:26:03 EDT 1998
_crashtime: struct {
    tv_sec = 746996332
    tv_usec = 145424
}
_boottime: struct {
    tv_sec = 746993148
    tv_usec = 92720
}
_config: struct {
    sysname = "OSF1"
    nodename = "system.dec.com"
    release = "V5.0"
    version = "331"
    machine = "alpha"
}
_cpu: 30
_system_string: 0xfffffc0000442fa8 = "AlphaServer 4100 5/400 4MB"
_avail_cpus: 1
_partial_dump: 1
_physmem(MBytes): 96
_panic_string: 0xfffffc000043cf70 = "kernel memory fault" [2]
_preserved_message_buffer_begin: [3]
struct {
    msg_magic = 0x63061
    msg_bufx = 0x56e
    msg_bufc = 0x432
    msg_bufc = "Alpha boot: available memory from 0x678000 to 0x6000000
Tru64 UNIX V5.0 (Rev. 1039); Tue Mar 30 08:26:03 EDT 1999
physical memory = 1024.00 megabytes.
available memory = 991.81 megabytes.
using 3924 buffers containing 30.65 megabytes of memory
tc0 at nexus
scc0 at tc0 slot 7
tcds0 at tc0 slot 6
asc0 at tcds0 slot 0
rz0 at asc0 bus 0 target 0 lun 0 (DEC RZ26 (C) DEC T384)
rz4 at asc0 bus 0 target 4 lun 0 (DEC RRD42 (C) DEC 4.5d)
tz5 at asc0 bus 0 target 5 lun 0 (DEC TLZ06 (C) DEC 0374)
```

```

ascl at tcds0 slot 1
rz8 at ascl bus 1 target 0 lun 0 (DEC      RZ57      (C) DEC 5000)
rz9 at ascl bus 1 target 1 lun 0 (DEC      RZ56      (C) DEC 0300)
fb0 at tc0 slot 8
  1280X1024
bba0 at tc0 slot 7
ln0: DEC LANCE Module Name: PMAD-BA
ln0 at tc0 slot 7
ln0: DEC LANCE Ethernet Interface, hardware address: 08-00-2b-2c-f3-83
Firmware revision: 5.1
PALcode: Tru64 UNIX version 1.21
AlphaServer 4100 5/400 4MB
lvm0: configured.
lvm1: configured.
<3>/var: file system full
<3>/var: file system full
<3>/var: file system full
<3>/var: file system full
<3>/var: file system full

trap: invalid memory ifetch access from kernel mode

      faulting virtual address:      0x0000000000000000
      pc of faulting instruction:    0x0000000000000000
      ra contents at time of fault:  0xffffffff000028951c
      sp contents at time of fault:  0xffffffff96199a48

panic: kernel memory fault
syncing disks... done
"
}
_preserved_message_buffer_end:
_kernel_process_status_begin: [4]
  PID COMM
00000 kernel idle
00001 init
00002 exception hdlr
00342 xdm
00012 update
00341 Xdec
00239 nfsiod
00113 syslogd
00115 binlogd
00240 nfsiod
00241 nfsiod
00340 csh
00124 routed
00188 portmap
00197 ypbind
00237 nfsiod
00249 sendmail
00294 internet_mom
00297 snmp_pe
00291 mold
00337 xdm
00325 lpd
00310 cron
00305 inetd
00489 tar
_kernel_process_status end:
_current_pid: 489 [5]
_current_tid: 0xffffffff863d36c0 [6]
_proc_thread_list_begin:
thread 0x863d36c0 stopped at [boot:1118,0xffffffff0000374a08] Source not available

```

```

_proc_thread_list_end:
_dump_begin: [7]
> 0 boot(reason = 0, howto = 0) ["../../../../src/kernel/arch/alpha/machdep.c":
1118, 0xfffffc0000374a08]
mp = 0xffffffff961962f8
nmp = 0xffffffff86333ab8
fsp = (nil)
rs = 5368785696
error = -1776721160
ind = 2424676
nbusy = 4643880

    1 panic(s = 0xfffffc000043cf70 = "kernel memory fault") ["../../../../src\
/kernel/bsd/subr_prf.c"\
:616, 0xfffffc000024ff60]
bootopt = 0

    2 trap() ["../../../../src/kernel/arch/alpha/trap.c":945, 0xfffffc0000381440]
t = 0xffffffff863d36c0
pcb = 0xffffffff96196000
task = 0xffffffff86306b80
p = 0xffffffff95aaf6a0
syst = struct {
    tv_sec = 0
    tv_usec = 0
}
nofault_save = 0
exc_type = 18446739675665756628
exc_code = 0
exc_subcode = 0
i = -2042898428
s = 2682484
ret = 536993792
map = 0xffffffff808fc5a0
prot = 5
cp = 0xffffffff95a607a0 =
i = 0
result = 18446744071932830456
pexcsum = 0xffffffff00000000
i = 16877
pexcsum = 0xffffffff00001000
i = 2682240
ticks = -1784281184
tv = 0xfffffc005000068

    3 _xentMM() ["../../../../src/kernel/arch/alpha/locore.s":949, 0xfffff\
c0000372dec]

_dump_end:

warning: Files compiled -g3: parameter values probably wrong
_kernel_thread_list_begin: [8]
thread 0x8632faf0 stopped at [thread_block:1427 ,0xfffffc00002ca3a0] Source\
not available
thread 0x8632f8d8 stopped at [thread_block:1427 ,0xfffffc00002ca3a0] Source\
not available
:
:

thread 0x8632d328 stopped at [thread_block:1400 +0x1c,0xfffffc00002ca2f8] \
Source not available
thread 0x8632d110 stopped at [thread_block:1400 +0x1c,0xfffffc00002ca2f8] \
Source not available
_kernel_thread_list_end:
_savedefp: 0xffffffff96199940 [9]

```

```

_kernel_memory_fault_data_begin: [10]
struct {
    fault_va = 0x0
    fault_pc = 0x0
    fault_ra = 0xfffffc000028951c
    fault_sp = 0xffffffff96199a48
    access = 0xffffffffffffffff
    status = 0x0
    cpunum = 0x0
    count = 0x1
    pcb = 0xffffffff96196000
    thread = 0xffffffff863d36c0
    task = 0xffffffff86306b80
    proc = 0xffffffff95aaf6a0
}
_kernel_memory_fault_data_end:
Invalid character in input
_uptime: .88 hours

_stack_trace_begin: [11]
> 0 boot(reason = 0, howto = 0) ["../../../../src/kernel/arch/alpha/machdep.c"\  
:1118, 0xfffffc0000374a08]
  1 panic(s = 0xfffffc000043cf70 = "kernel memory fault") ["../../../../src\  
/kernel/bsd/subr_prf.c":616, 0xfffffc000024ff60]
  2 trap() ["../../../../src/kernel/arch/alpha/trap.c":945, 0xfffffc0000381\  
440]
  3 _XentMM() ["../../../../src/kernel/arch/alpha/locore.s":949, 0xfffffc000\  
0372dec]
_stack_trace_end:
_savedefp_exception_frame_(savedefp/33X): [12]
ffffffff96199940: 0000000000000000 fffffc000046f888
ffffffff96199950: ffffffff863d36c0 0000000079c2c93f
ffffffff96199960: 000000000000007d 0000000000000001
ffffffff96199970: 0000000000000000 fffffc000046f4e0
ffffffff96199980: 0000000000000000 ffffffff961962f8
ffffffff96199990: 0000000140012b20 0000000000000000
ffffffff961999a0: 0000000140045690 0000000000000000
ffffffff961999b0: 00000001400075e8 0000000140026240
ffffffff961999c0: ffffffff961999af0 ffffffff8635adc0
ffffffff961999d0: ffffffff961999ac0 00000000000001b0
ffffffff961999e0: fffffc00004941b8 0000000000000000
ffffffff961999f0: 0000000000000001 fffffc000028951c
ffffffff961999a0: 0000000000000000 0000000000000fff
ffffffff961999a10: 0000000140026240 0000000000000000
ffffffff961999a20: 0000000000000000 fffffc000047acd0
:
:
ffffffff961999a30: 0000000000901402 0000000000001001
ffffffff961999a40: 0000000000002000
_savedefp_exception_frame_ptr: 0xffffffff96199940
_savedefp_stack_pointer: 0x140026240
_savedefp_processor_status: 0x0
_savedefp_return_address: 0xfffffc000028951c
_savedefp_pc: 0x0
_savedefp_pc/i:

can't read from process (address 0x0)
_savedefp_return_address/i:
[spec_open:997, 0xfffffc000028951c] bis r0, r0, r19
_kernel_memory_fault_data.fault_pc/i:

can't read from process (address 0x0)
_kernel_memory_fault_data.fault_ra/i:
[spec_open:997, 0xfffffc000028951c] bis r0, r0, r19

```

```

_kdbx_sum_start:
Hostname : system.dec.com
cpu: AlphaServer 4100 5/400      avail: 1
Boot-time:      Tues Jul  7 10:33:25 1998
Time:   Mon Jul 13 13:58:52 1998
Kernel : OSF1 release V5.0 version 688.2 (alpha)
_kdbx_sum_end:
_kdbx_swap_start: 13
      Swap device name          Size      In Use      Free
-----
/dev/rz0b                    131072k    10560k    120512k  Dumpdev
                              16384p      1320p     15064p
-----
Total swap partitions:  1      131072k    10560k    120512k
                              16384p      1320p     15064p
_kdbx_swap_end:
_kdbx_proc_start: 14
Addr      PID  PPID  PGRP  UID  NICE  SIGCATCH  P_SIG      Event  Flags
-----
v0x95aaf6a0 489  340  489   0    0    00000000  00000000  NULL   in pagv ctty
v0x95aad5d0 342  337  342   0    0    00000000  00000000  NULL   in pagv ctty
v0x95aad8f0 341  337  341   0    0    00000000  00000000  NULL   in pagv
:
:
v0x95aad2b0  1    0    1    0    0    00000000  00000000  NULL   in omask pagv
v0x95aad120  0    0    0    0    0    00000000  00000000  NULL   in sys
kdbx_proc_end:

Audit subsystem not installed
#
_crash_data_collection_finished:

```

- 1 The first several lines of output display the contents of system variables that give statistics about the crash, such as:
 - The kernel image file and crash core file from which `crashdc` collected data.
 - The operating system version.
 - The time of the crash and the time at which the system was rebooted.
 - Whether data is from a partial or full dump. (Data is from a partial dump when the value of the `partial_dump` variable is 1. Data is from a full dump when the value of this variable is 0.)
 - The platform on which the operating system is running; an AlphaServer 4100 in this case.
 - The amount of physical memory available on the system.
- 2 The `_panic_string` label marks the message that indicates why the crash occurred. In this case the message is kernel memory fault, indicating that a memory operation failed in the kernel.
- 3 The preserved message buffer contains status and other information about the devices connected to the system: Notice the following message:

```
trap:  invalid memory ifetch access from kernel mode
```

This message describes the kernel memory fault and indicates that the kernel was unable to fetch a needed instruction.

The preserved message buffer also contains the faulting virtual address, the pc of the instruction that failed, the contents of the return address register, and the stack pointer at the time of the memory fault.

- 4 The kernel process status list shows the processes that were active at the time of the crash.
- 5 The `_current_pid` label marks the process ID of the process that was executing at the time of the crash. In this case, it is the `tar` process, which is identified as process 489 in the kernel process status list.
- 6 The `_current_tid` label marks the address of the thread that was executing at the time of the crash.
- 7 The dump section shows information about the variables passed to the routines executing at the time of the crash. In this case, the dump displays variable information for the `boot`, `panic`, and `trap` functions.
- 8 The kernel thread list shows the threads of execution in the kernel. This information can be helpful for verifying which routine called the `panic` function.
- 9 The `savedefp` variable contains a pointer to the exception frame.
- 10 The kernel memory fault data displays the following information, recorded at the time of the memory fault:
 - The `fault_va` variable contains the faulting virtual address.
 - The `fault_pc` variable contains the pc.
 - The `fault_ra` variable contains the return address of the calling routine.
 - The `fault_sp` variable contains the stack pointer.
 - The `access` variable contains the access code, which is zero (0) for read access, 1 for write access, and -1 for execute access.
 - The `status` variable contains the process status register.
 - The `cpunum` variable contains the number of the CPU that faulted.
 - The `count` variable contains the number of CPUs on the system.
 - The `pcb` variable contains a pointer to the process control block.
 - The `thread` variable contains a pointer to the current thread.
 - The `task` variable contains a pointer to the current task.
 - The `proc` variable contains the address of the process status table.

- [11]** The `_stack_trace_begin` line begins a trace of the current thread block's stack at the time of the crash. In this case the `_xentMM` function called the `trap` function. The `trap` function called the `panic` function, which called the `boot` function and the system crashed.
- [12]** The exception frame is a stack frame created to store the state of the process running at the time of the exception. It stores the registers and `pc` associated with the process. To determine where registers are stored in the exception frame, refer to the `/usr/include/machine/reg.h` header file.
- [13]** Swap information is shown to help you determine whether swap space is sufficient.
- [14]** The process table gives information about the processes active at the time of the crash. The information includes:
 - The process ID of each process.
 - The process ID of the parent process for each process.
 - The process group ID for each process.
 - The UID of the user that started each process. In this case all processes are started by `root`.
 - The priority at which the process was running at the time of the memory fault.
 - The event the process was waiting for, if any. An event might be the completion of an input or output request, for example.
 - Any flags assigned to the process. For example, the `ctty` flag indicates that the process has a controlling terminal and, the `sys` flag indicates that the process is a swapper or pager process.

Index

A

abscallout kdbx extension, 2–19
access variable, A–6
addr_to_proc function, 3–3
alias command, 2–13
Alpha hardware architecture
 documentation, 1–1
arp kdbx extension, 2–16
array
 using in a kdbx extension, 3–27e
 using in kdbx extension, 3–28e
array_action kdbx extension, 2–16
array_element function, 3–4
array_element_val function, 3–4
array_size function, 3–6

B

boot function, 1–5
bootstrap-linked kernel
 debugging, 1–1
breakpoint
 setting on an SMP system, 2–44
buf kdbx extension, 2–18
build system, 2–38

C

call stack
 of user program, examining in crash dump, 2–8
callout kdbx extension, 2–18
cast function, 3–6
cast kdbx extension, 2–19
cc command

 using to compile a kdbx extension, 3–35
check_args function, 3–7
check_fields function, 3–7
complex lock
 displaying debug information for, 2–11
config kdbx extension, 2–19
context command, 2–14
context function, 3–8
convert kdbx extension, 2–20
coredata command, 2–14
count variable, A–6
cpunum variable, A–6
cpustat extension, 2–20
crash data collection, 2–44, A–1
crash dump analysis, 1–1
 collecting data with crashdc, 2–44
 examples of, 4–1
 for SMP systems, 4–9
 guidelines for, 4–1
 specifying location of loadable modules for, 2–4
 viewing user program stack, 2–8
crash dump file
 analyzing, 1–5
 example of using dbx to examine, 4–2
 example of using kdbx to examine, 4–3
 guidelines for analyzing, 4–1
 invoking dbx debugger to examine, 2–2
 invoking kdbx debugger to examine, 2–12

crash-data.n file
 explanation of contents, A-1
crashdc command
 explanation of output from, A-1
crashdc utility, 2-44
customizing kdbx debugger environment, 2-13

D

data structure
 displaying format of with dbx debugger, 2-6
 displaying with dbx debugger, 2-6
data types used by kdbx extensions, 3-2
DataStruct data type, 3-3
dbx command, 2-14
dbx debugger, 2-2
 breakpoint handling on an SMP system, 2-44
 debugging kdbx extensions with, 3-36
 debugging kernel threads with, 4-8
 displaying call stack of user program after kernel crash with, 2-8
 displaying format of data structures with, 2-6
 displaying preserved message buffer, 2-10
 displaying variable and data structure with, 2-6
 examining exception frames with, 2-7
 example of using for crash dump analysis, 4-2, 4-4
 example of using for identifying hardware exception, 4-4
 identifying cause of crash on SMP system, 4-9
 kernel debugging flag, 2-2
 syntax for address formats, 2-2

 syntax for examining dump files, 2-2
 using dbx commands in kdbx extension, 3-9
dbx function, 3-9
debugging kernel threads with dbx, 4-8
debugging kernels
 (*See* kernel debugging)
debugging tools
 crashdc utility, 2-44
 dbx debugger, 2-2
 kdbx debugger, 2-12
 kdebug debugger, 2-37
deref_pointer function, 3-9
device configuration
 displaying, 2-10
dis kdbx extension, 2-21
disassembling instructions, 2-21
disassembling return addresses, 4-6
disassembling the pc value, 4-6
dump file
 (*See* crash dump file)
dump function, 1-5

E

exception frame, A-7
 examining with dbx debugger, 2-7
export kdbx extension, 2-21
extensions to kdbx debugger, 2-15

F

fault_pc variable, A-6
fault_ra variable, A-6
fault_sp variable, A-6
fault_va variable, A-6
field_errors function, 3-10
FieldRec data type, 3-3
file command

using to determine type of kernel,
1-2

file kdbx extension, 2-21

firmware version

displaying, 2-10

format_addr function, 3-10

free_sym function, 3-11

G

gateway system, 2-38

global symbols

using in kdbx extension, 3-34e

H

hardware exception

example of debugging, 4-4

help command, 2-14

I

inpcb kdbx extension, 2-22

instructions

disassembling using kdbx, 2-21

K

kdbx debugger, 2-12, 2-15, 3-2

(*See also* specific library
routines; specific kdbx
extensions)

breakpoint handling on an SMP
system, 2-44

command aliases, 2-15

command syntax, 2-12

commands, 2-13

compiling custom extensions to,
3-35

customizing environment of, 2-13

debugging extensions to, 3-36

example of using for crash dump
analysis, 4-3, 4-7

example of using for identifying

hardware exception, 4-4

executing extensions to, 2-14

extensions to, 3-22

initialization files, 2-13

library functions for extensions to,
3-2

special data types, 3-2

using extensions to, 2-15

writing extensions for

using arrays, 3-28e

using arrays template, 3-27e

using global symbols, 3-34e

using linked lists, 3-24e

using lists template, 3-23e

writing extensions to, 3-1

kdbx extensions

checking arguments passed to, 3-7

compiling, 3-35

library routines for writing, 3-1

using arrays, 3-28e

using global symbols, 3-34e

using linked lists, 3-24e

using lists template, 3-23e

kdbxrc file, 2-13

kdebug debugger, 1-3, 2-37

invoking, 2-41

problems with setup of, 2-42

requirements for, 2-38

setting up, 2-39

kernel

determining boot method of, 1-2

kernel crash

displaying call stack of user
program after, 2-8

kernel program

debugging, 1-3

kernel thread list

location of in crashdc output, A-6

kps command, 4-6

krash function, 3–11

L

ld command

using to build a kernel image file,
1–2

libkdbx.a library, 3–1

library functions

for extensions to kdbx debugger,
3–2

library routines

for writing kdbx extensions, 3–2

linked list

using in a kdbx extension, 3–24e

list_action kdbx extension, 2–22

list_nth_cell function, 3–13

loadable modules

specifying location of for crash
dumps, 2–4

lock

(*See* complex lock, simple lock)

lockinfo kdbx extension, 2–25

lockmode system attribute, 2–10

lockstats kdbx extension, 2–24

M

machine_slot structure, 4–11

message buffer, preserved

examining with dbx debugger, 2–10

modules

loadable, specifying location of for
crash dumps, 2–4

mount kdbx extension, 2–26

N

namecache kdbx extension, 2–27

new_proc function, 3–14

next_number function, 3–14

next_token function, 3–15

O

ofile kdbx extension, 2–27

operating system version

displaying, 2–10

location of in crashdc output, A–5

P

p command, 4–3

paddr kdbx extension, 2–28

panic function, 1–5

panic string

location of in crashdc output, A–5

where stored, 4–1

paniccpu variable, 4–9

panicstr variable

example of displaying, 4–3

pc value

determining with kdbx, 4–6

disassembling, 4–6

pcb kdbx extension, 2–28

pcb variable, A–6

PID

displaying, 4–6

pointer

casting to a data structure, 3–6

pr command, 2–14

preserved message buffer

contents of, A–5

examining with dbx debugger, 2–10

example of displaying, 4–6

print command, 2–15

print function, 3–16

print_status function, 3–16

printf kdbx extension, 2–28

proc kdbx extension, 2–29

procaddr kdbx extension, 2–30

process control block

displaying for a thread, 2–28

process ID

location of in crashdc output, A–6

process table, A–7

displaying, 2–29

Q

quit command, 2–15
quit function, 3–17

R

read_field_vals function, 3–17
read_line function, 3–18
read_memory function, 3–18
read_response function, 3–19
read_sym function, 3–20
read_sym_addr function, 3–20
read_sym_val function, 3–21
reg.h header file
(See /usr/include/machine/reg.h
header file)
remote debugging, 2–37
**requirements for kdebug
debugger**, 2–38

S

savedefp variable, 2–8
location of in crashdc output, A–6
setting up the kdebug debugger,
2–39
simple lock
displaying debug information for,
2–11
sizer command
using to determine type of kernel,
1–2
slock_debug array, 2–11
SMP system
debugging on, 2–10
determining on which CPU a panic
occurred, 4–9
socket kdbx extension, 2–30
software panic
example of debugging, 4–2
source command, 2–15

stack of user program

examining in crash dump, 2–8

stack trace

example of, 4–6
multiple panic messages in, 4–11

Status data type, 3–2

status variable, A–6

StatusType data type, 3–2

struct_addr function, 3–21

sum command, 4–4

sum kdbx extension, 2–30

swap kdbx extension, 2–31

swap space

displaying with kdbx, 2–31

sysconfig command

using to set the lockmode attribute,
2–10

system

displaying information about with
kdbx, 2–30

System boot method

determining, 1–2

system crash

identifying the cause of, 4–2
process of, 1–5
reasons for, 1–5
using crashdc command to collect
data from, 2–44
using dbx to find the cause of, 2–2
using kdbx to find the cause of,
2–12

system.kdbxrc file, 2–13

T

t command, 4–4

task ID

location of in crashdc output, A–6

task kdbx extension, 2–31

task variable, A–6

tcb table

- displaying using the inpcb kdbx extension, 2-22
- test system**, 2-38
- testing kernel programs**, 2-37
- thread**
 - displaying the process control block for, 2-28
- thread kdbx extension**, 2-32
- thread variable**, A-6
- to_number function**, 3-22
- trace command**, 4-3
- trace kdbx extension**, 2-32
- tracing execution**
 - during crash dump analysis, 4-4
 - on an SMP system, 4-11
- tset command**, 4-11

U

- u kdbx extension**, 2-33

- ucred kdbx extension**, 2-34
- udb table**
 - displaying using the inpcb kdbx extension, 2-22
- unalias command**, 2-15
- unaliasall kdbx extension**, 2-36
- user program**
 - examining call stack of in crash dump, 2-8
- /usr/include/machine/reg.h header file**, 2-8
- ustname structure**
 - example of displaying, 4-6

V

- variable**
 - displaying with dbx debugger, 2-6
- vnode extension**, 2-36